Veeam Plug-in for Proxmox VE

Version 3

User Guide

November, 2025

NOTE

Read the End User Software License Agreement before using the accompanying software programs. Using any part of the software indicates that you accept the terms of the End User Software License Agreement.

Contents

Contacting Veeam Software

At Veeam Software we value feedback from our customers. It is important not only to help you quickly with your technical issues, but it is our mission to listen to your input and build products that incorporate your suggestions.

Customer Support

Should you have a technical concern, suggestion or question, visit the Veeam Customer Support Portal to open a case, search our knowledge base, reference documentation, manage your license or obtain the latest product release.

Company Contacts

For the most up-to-date information about company contacts and office locations, visit the Veeam Contacts Webpage.

Online Support

If you have any questions about Veeam products, you can use the following resources:

- Full documentation set: veeam.com/documentation-guides-datasheets.html
- Veeam R&D Forums: forums.veeam.com

Overview

Veeam Plug-in for Proxmox VE is a software component developed for protection and disaster recovery tasks for Proxmox Virtual Environment (Proxmox VE). This component comes as part of the Veeam Backup & Replication solution and allows you to perform the following operations:

- Create backups of Proxmox VE VMs and store them in backup repositories.
- Create VeeamZIP backups of Proxmox VE VMs.
- Create several instances (copies) of the same backup data in different locations.
- Restore VMs from Proxmox VE VM backups to the Proxmox VE environment.
- Restore VMs from VMware ESXi and Microsoft Hyper-V to the Proxmox VE environment.
- Restore VMs from Nutanix AHV and oVirt KVM backups to the Proxmox VE environment.
- Restore VMs from Microsoft Azure, Amazon Web Services (AWS) and Google Cloud backups to the Proxmox VE environment.
- Restore physical machines from backups created by Veeam Agents to the Proxmox VE environment.
- Restore VMs from Proxmox VE backups to Microsoft Azure, Amazon Web Services (AWS) and Google Cloud environments.
- Restore VMs from Proxmox VE backups to Nutanix AHV, oVirt KVM, Scale Computing HyperCore environments.
- Perform Instant Recovery of Proxmox VE VMs to VMware vSphere and Microsoft Hyper-V environments.
- Restore application items (such as Microsoft Active Directory, Microsoft Exchange, Microsoft SharePoint, Microsoft DNS, PostgreSQL, Oracle Database and Microsoft SQL Server).
- Restore files and folders of Proxmox VE VM guest OSes.
- Export disks of backed-up Proxmox VE VMs to VMDK, VHD and VHDX formats.
- Mount disks of backed-up Proxmox VE VMs to any server and access data in the read-only mode.

Solution Architecture

Since Veeam Plug-in for Proxmox VE is integrated with Veeam Backup & Replication, the solution architecture comprises the following set of components:

- Proxmox VE server
- Backup server
- Veeam Plug-in for Proxmox VE
- Backup repositories
- Workers

Proxmox VE Server

A Proxmox VE server is standalone host or cluster node that runs the Proxmox VE software. Veeam Plug-in for Proxmox VE uses the server to access such Proxmox VE resources as storage, networks and VMs while performing backup and restore operations.

Backup Server

A backup server is an either Windows-based or Linux-based physical or virtual machine on which Veeam Backup & Replication is installed. The backup server is the configuration, administration and management core of the backup infrastructure. It coordinates backup and restore operations, controls job scheduling and manages resource allocation.

Veeam Plug-in for Proxmox VE

Veeam Plug-in for Proxmox VE is an architecture component that enables integration between the backup server and other components of the backup infrastructure. Veeam Plug-in for Proxmox VE allows Veeam Backup & Replication to connect to the Proxmox VE server, and to perform data protection and disaster recovery tasks with Proxmox VE resources.

Backup Repositories

A backup repository is a storage location where Veeam Backup & Replication stores backups of protected Proxmox VE VMs.

To communicate with backup repositories, Veeam Backup & Replication uses Veeam Data Mover — the service that is responsible for data processing and transfer. By default, Veeam Data Mover runs on the repositories themselves. If a repository cannot host Veeam Data Mover, it starts on a gateway server — a dedicated component that "bridges" the backup server and workers. For more information, see the Veeam Backup & Replication User Guide, section Gateway Server.

Workers

A worker is a Linux-based VM that resides on the Proxmox VE host and processes backup workloads when transferring data to and from backup repositories.

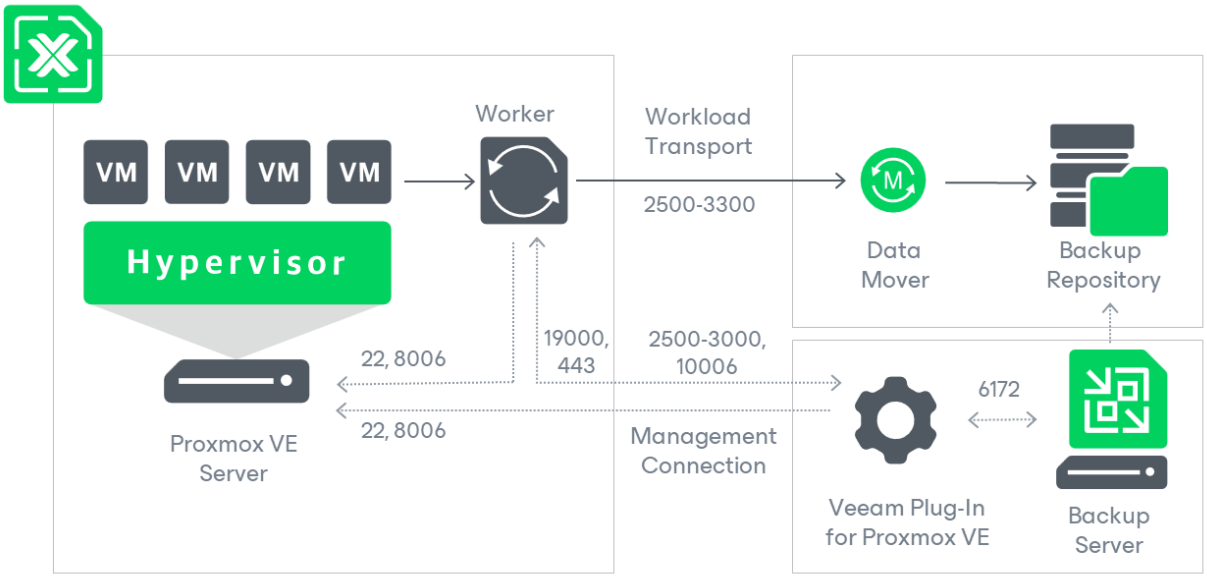

VM Backup

To produce backups of VMs, Veeam Backup & Replication runs backup jobs. A backup job is a collection of settings that define the way backup operations are performed: what data to back up, where to store backups, when to start the backup process, and so on.

Veeam Backup & Replication does not install agent software inside VMs to back up VM data — it uses native Proxmox VE capabilities instead. During every backup session, Veeam Backup & Replication creates a Proxmox VE copy-on-write snapshot of each VM added to a backup job. The snapshot is further used to create a VM backup.

How to Protect VMs

1. Check system requirements and account permissions.
2. Add backup repositories.
3. Connect the Proxmox VE server.
4. Configure worker settings.
5. Configure email settings and notifications.
6. Complete the New Backup Job wizard.

How VM Backup Works

Veeam Backup & Replication performs VM backup in the following way:

1. Launches a worker on the same host where the processed VM resides.

   If no worker is deployed on the host, Veeam Backup & Replication launches a worker that is deployed on any other Proxmox VE host connected to the backup infrastructure.

2. Connects to the Proxmox VE server and creates a copy-on-write snapshot of the processed VM.
3. Uses the worker to read data from disks that are attached to the processed VM, compares it to the data written to the snapshot created at the step 2, excludes the changes and transfers the resulting data to the target repository — and stores it in the native Veeam format.

   To reduce the amount of data read from VM disks, Veeam Backup & Replication uses the changed block tracking (CBT) mechanism: during incremental backup sessions, Veeam Backup & Replicationcompares the current disk content with the backed-up content and reads only those data blocks that have changed since the previous backup session. If CBT cannot be used, Veeam Backup & Replication reads all data from the VM disks. For more information, see Changed Block Tracking.

   Veeam Backup & Replication compresses and deduplicates data saved to repositories.

4. Removes the created snapshot and suspends the worker when the backup session completes.

Backup Chain

Veeam Backup & Replication creates a new backup file in a backup repository during every backup session. A sequence of backup files created during a set of backup sessions makes up a backup chain. Each backup chain contains data for one VM only. If a backup job includes several VMs, Veeam Backup & Replication creates one backup chain for each VM processed by the job.

The backup chain includes backup files of the following types:

- VBK — a full backup file stores a copy of the full VM image.
- VIB — incremental backup files store incremental changes of the VM image.
- VBM — backup metadata files store information about the backup job, VMs processed by the backup job, number and structure of backup files, restore points, and so on. Metadata files facilitate import of backups, backup mapping and other operations.

Full and incremental backup files act as restore points for backed-up VMs that let you roll back VM data to the necessary state. To recover a VM to a specific point in time, the chain of backup files created for the VM must contain a full backup file and a set of incremental backup files dependent on the full backup file.

If some file in the backup chain is missing, you will not be able to roll back to the necessary state. For this reason, you must not delete individual backup files from the backup repository manually. Instead, you must specify retention policy settings that will let you maintain the necessary number of backup files in the backup repository. For more information, see Backup Retention.

Changed Block Tracking

The changed block tracking (CBT) mechanism allows Veeam Backup & Replication to reduce the amount of data read from processed VMs, and to increase the speed and efficiency of incremental backups:

- During a full backup session Veeam Backup & Replication reads only written data blocks, while unallocated data blocks are filtered out.
- During an incremental backup session, Veeam Backup & Replication reads only those data blocks that have changed since the previous backup session.

To detect unallocated and changed data blocks, CBT relies on the QEMU Dirty Bitmaps functionality:

1. During the first (full) backup session, Veeam Backup & Replication creates a bitmap for each disk that is attached to a processed VM.
2. During subsequent sessions, Veeam Backup & Replication uses the created bitmaps to compare the contents of disks backed up during the previous backup session and the current disk contents. This allows Veeam Backup & Replication to detect data blocks that have changed since the previous backup session. As soon as a new backup is created, Veeam Backup & Replication updates the bitmaps to include the latest changes.

Limitations for Changed Block Tracking

Due to Proxmox VE technical limitations, bitmaps created for disks in the RAW and VMDK formats are automatically removed as soon as VMs that have these disks attached are powered off or restarted. Therefore, Veeam Backup & Replication may not be able to use CBT when processing those VMs and trying to detect data blocks that have changed since the previous backup session. If CBT cannot be used, Veeam Backup & Replication reads the whole content of VM disks and compares it with backed-up data that already exists in backup repositories. In this case, the completion time of incremental backups may occur to grow.

> **NOTE**
>
> This limitation does not apply to disks in the QCOW2 format.

Backup Retention

Veeam Backup & Replication retains the number of latest restore points defined in job scheduling settings as described in section Creating Backup Jobs. For backup chains created by jobs without scheduled active or synthetic full backups, Veeam Backup & Replication applies forever forward incremental backup retention policy. For backup chains created by jobs that regularly produce active or synthetic full backups, Veeam Backup & Replication applies forward incremental backup retention policy.

> **NOTE**
>
> For backup chains created by jobs that no longer exist, Veeam Backup & Replication applies a separate retention mechanism as described in the Veeam Backup & Replication User Guide, section Background Retention.

Forever Forward Incremental Backup Retention Policy

To track and remove redundant restore points from a forever forward incremental backup chain, Veeam Backup & Replication performs the following actions once a day:

1. Veeam Backup & Replication checks the configuration database to detect backup chains with restore points that are older than the specified time limit.
2. If a redundant restore point exists in a backup chain, Veeam Backup & Replication transforms the backup chain in the following way:
   a. Rebuilds the full backup to include there data of the incremental backup that follows the full backup. To do that, Veeam Backup & Replication injects into the full backup data blocks from the earliest incremental backup in the chain. This way, the full backup 'moves' forward in the standard backup chain.

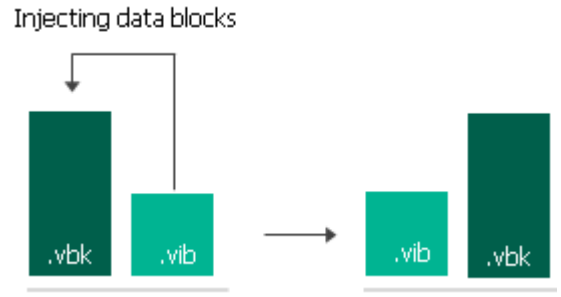

   b. Removes the earliest incremental backup from the chain as redundant — this data has already been injected into the full backup.

3. Veeam Backup & Replication repeats step 2 for all other redundant restore points found in the backup chain until all the restore points are removed. As data from multiple restore points is injected into the rebuilt full backup, Veeam Backup & Replication ensures that the backup chain is not broken and that you will be able to recover your data when needed.

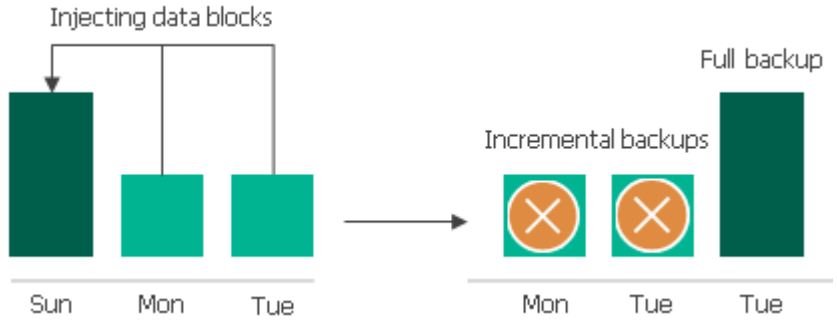

Forward Incremental Backup Retention Policy

To track and remove redundant restore points from a forward incremental backup chain, Veeam Backup & Replication performs the following actions once a day:

1. Veeam Backup & Replication checks the configuration database to detect forward incremental backup chains where a new full backup has been created (which starts a new backup chain fragment).
2. Veeam Backup & Replication checks the following whether the period to keep restore points in the new chain fragment has reached the allowed time limit.
3. If the new backup chain fragment has reached the limit of allowed restore points, Veeam Backup & Replication removes all restore points of the older backup chain fragment.

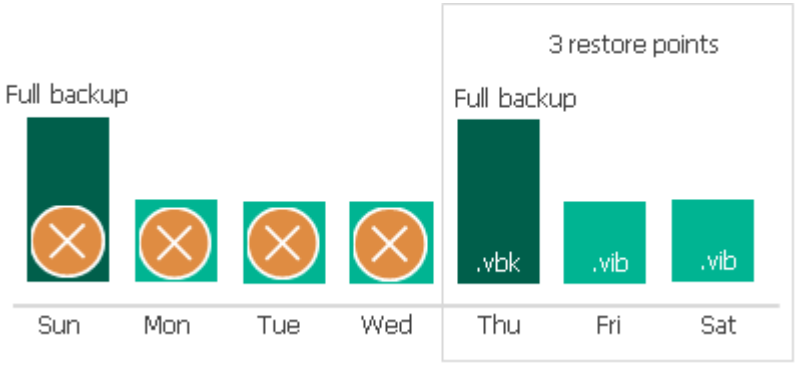

Backup Methods

Veeam Backup & Replication provides the following methods for creating backup chains:

- **Forever forward incremental**

  When the forever forward incremental backup method is used, Veeam Backup & Replication creates a backup chain that consists of the first full backup file (VBK) and a set of forward incremental backup files (VIBs) following it. For more information, see Forever Forward Incremental Backup.

  This backup method helps you save space on the backup storage because Veeam Backup & Replication stores only one full backup file and removes incremental backup files once the retention period is exceeded.

- **Forward incremental**

  When the forward incremental backup method is used, Veeam Backup & Replication creates a backup chain that consists of multiple full backup files (VBKs) and sets of forward incremental backup files (VIBs) following each full backup file. Full backups created using the synthetic full or active full method split the backup chain into shorter series. This lowers the chances of losing the backup chain completely and makes this backup method the most reliable. For more information, see Forward Incremental Backup.

  This backup method requires more storage space than other methods because the backup chains contains multiple full backup files and sometimes Veeam Backup & Replication stores more restore points than specified in the retention policy settings due to the specifics of the forward incremental retention policy.

Forever Forward Incremental Backup

To create a backup chain for a VM protected by a backup job that is not configured to produce full backups, Veeam Backup & Replication implements the forever forward incremental backup:

1. During the first (full) backup session, Veeam Backup & Replication copies the full VM image and creates a full backup file in the backup repository. The full backup file becomes a starting point in the backup chain.
2. During subsequent backup sessions, Veeam Backup & Replication copies only those data blocks that have changed since the previous backup session, and stores these data blocks to incremental backup files in the backup repository. The content of each incremental backup file depends on the content of the full backup file and the preceding incremental backup files in the backup chain.

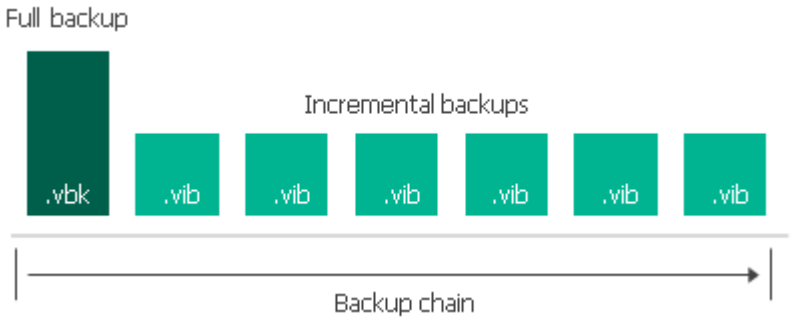

Forward Incremental Backup

To create a backup chain for a VM protected by a backup job that is configured to produce full backups, Veeam Backup & Replication implements the forward incremental backup method:

1. During the first (full) backup session, Veeam Backup & Replication copies the full VM image and creates a full backup file in the backup repository. The full backup file becomes a starting point in the backup chain.

2. During subsequent backup sessions, Veeam Backup & Replication copies only those data blocks that have changed since the previous backup session, and stores these data blocks to incremental backup files in the backup repository. The content of each incremental backup file depends on the content of the full backup file and the preceding incremental backup files in the backup chain.

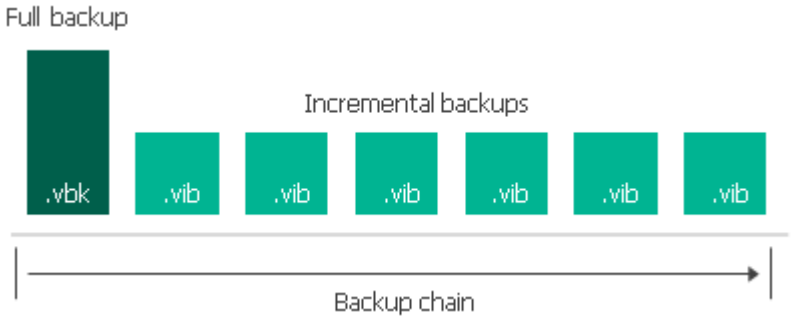

3. On a day when the synthetic full or active full backup is scheduled, Veeam Backup & Replication creates a full backup file and adds it to the backup chain. Incremental restore points produced after this full backup file use it as a new starting point.

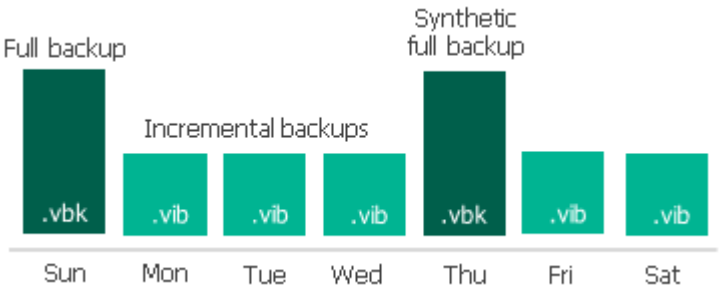

Active Full Backup

In some cases, you need to regularly create a full backup. For example, your corporate backup policy may require that you create a full backup on weekend and run incremental backup on work days. To let you conform to these requirements, Veeam Backup & Replication allows you to create active full backups (either manually or automatically according to a specific schedule).

When creating an active full backup, Veeam Backup & Replication starts a new backup chain for the VM. All further created incremental backups use the latest active full backup file as a new starting point. The old full backup file from the old backup chain remains on disk until it is automatically deleted according to the retention policy.

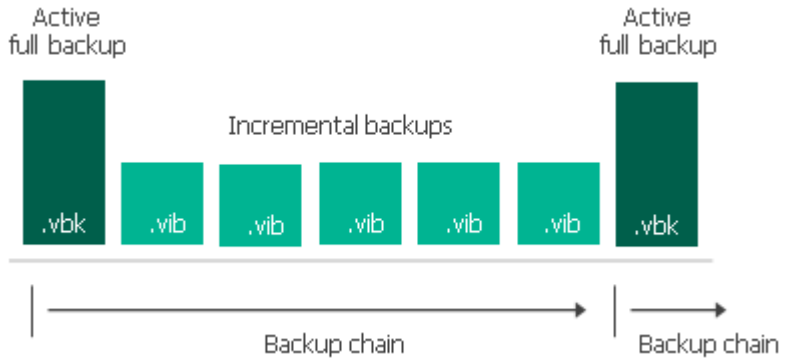

The active full backup session starts at the same time when the backup job is scheduled. For example, if you schedule the backup job to run at 12:00 AM Sunday through Friday, and schedule active full backup to be created on Saturday, Veeam Backup & Replication will start a backup job session that will produce an active full backup at 12:00 AM on Saturday.

If the backup job is not scheduled to run automatically or is disabled, Veeam Backup & Replication will not perform active full backup. If a regular backup session and an active full backup session are scheduled on the same day, Veeam Backup & Replication will produce an active full backup — an incremental backup that should have been created by the regular backup session will not be added to the backup chain. However, if you run the backup job again on the same day manually, Veeam Backup & Replication will perform incremental backup in a regular manner.

Synthetic Full Backup

In some situations, running active full backups periodically may not be an option. Active full backups are resource-intensive and consume considerable amount of network bandwidth. As an alternative, you can create synthetic full backups that also produce VBK files and contain data of the whole VM. However, while creating synthetic full backups, Veeam Backup & Replication does not retrieve VM data from the cluster but processes the data that is already stored in the backup repository.

To create a synthetic full backup, Veeam Backup & Replication performs the following operations:

1. Veeam Backup & Replication creates a regular incremental backup and adds it to the backup chain.

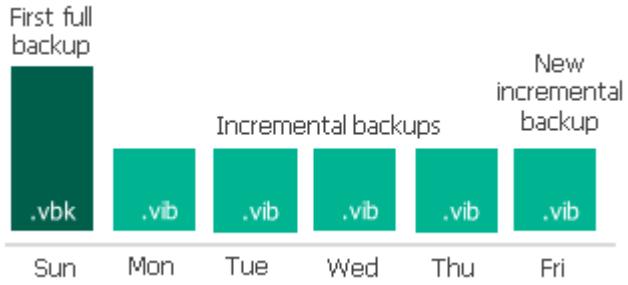

2. Veeam Backup & Replication creates a new synthetic full backup using backup files that are already available in the backup chain, including the newly created incremental backup file.

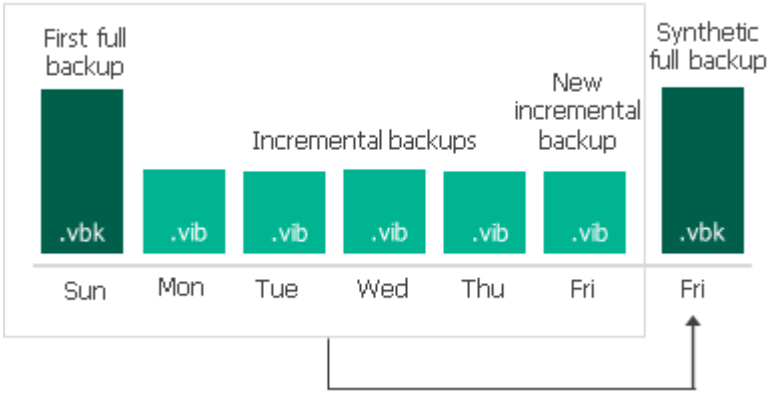

3. Veeam Backup & Replication deletes the created incremental backup as its data is already incorporated in the synthetic full backup.

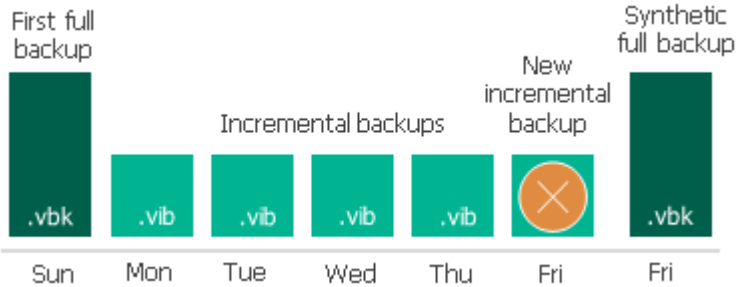

When creating a synthetic full backup, Veeam Backup & Replication starts a new backup chain for the VM. All further created incremental backups use the latest full backup file as a new starting point. The old full backup file from the old backup chain remains on disk until it is automatically deleted according to the retention policy.

> **NOTE**
>
> The synthetic full backup session starts only on the day when the backup job is scheduled. For example, if you schedule the backup job to run at 12:00 AM Sunday through Friday, and schedule synthetic full backup to be created on Saturday, Veeam Backup & Replication will never start a backup job session that will produce a synthetic full backup.

If the backup job is not scheduled to run automatically or is disabled, Veeam Backup & Replication will not perform synthetic full backup. If a regular backup session and a synthetic full backup session are scheduled on the same day, Veeam Backup & Replication will produce a synthetic full backup — an incremental backup that should have been created by the regular backup session will not be added to the backup chain. However, if you run the backup job again on the same day manually, Veeam Backup & Replication will perform incremental backup in a regular manner.

VM Restore

Veeam Backup & Replication offers the following restore options:

- Entire VM Restore — restores an entire VM from a backup. You can restore one or more VMs at a time, to the original location or to a new location.
- File-level recovery — recovers individual VM files and folders from a backup. You can download the necessary files and folders to a local machine, or restore the files and folders of the source VM to the original location.

You can restore VM data to the most recent state or to any available restore point.

Entire VM Restore

To restore a VM, Veeam Backup & Replication performs the following steps:

1. [This step applies only if you perform restore to the original location and if the source VM is still present in the location] Connects to the Proxmox VE server over REST API to power off and remove the source VM.
2. Launches a worker on same host where the processed VM resides.

   If no worker is deployed on the host, Veeam Backup & Replication launches a worker that is deployed on any other Proxmox VE host connected to the backup infrastructure.
3. Connects to the Proxmox VE server over REST API, configures a VM and creates empty virtual disks in the target location.

   The number of empty disks equals the number of disks attached to the backed-up VM.
4. Restores backed-up data to the empty disks and restores them to the configured VM.

   If multiple disks are attached to the backed-up VM, these disks are restored sequentially, one disk at a time.
5. Suspends the worker when the restore session completes.

> **NOTE**
>
> If multiple VMs are added to the restore session, these VMs are processed in parallel.

To learn how to restore an entire VM, see Performing VM Restore.

File-Level Recovery

To recover VM files and folders from a backup, Veeam Backup & Replication performs the following steps:

1. Mounts disks of the backed-up VM to the mount host specified for the recovery operation.
2. Launches the Veeam Backup Browser.

   The Veeam Backup Browser displays the directory structure of the backed-up VM. In the browser, you select the necessary files and folders to restore.
3. Restores the selected files and folders to the original location or to a new location.
4. Detaches the disks from the mount host.

To learn how to recover individual VM files and folders, see Performing File-Level Restore.

Retention Policies

Image-level backups created by jobs are not kept forever — they are removed according to retention policy settings specified while creating the jobs as described in section Creating Backup Jobs.

Restore points in the backup chain are stored only for the allowed period of time (in days). If a restore point is older than the specified time limit, Veeam Backup & Replication removes it from the backup chain. To learn how Veeam Backup & Replication applies retention policies to forever forward incremental and forward incremental backup chains, see Backup Retention.

Planning and Preparation

Before you start using Veeam Plug-in for Proxmox VE, check supported virtualization platforms, system requirements, limitations, permissions and network ports used for data transmission.

System Requirements

Before you start using Veeam Plug-in for Proxmox VE, make sure the virtual environment and the backup infrastructure components meet the following requirements.

Specification	Requirement
Hypervisor	Kernel-based Virtual Machine (KVM) must be installed on x86 hardware that supports virtualization capabilities.
Virtualization Platform	Veeam Plug-in for Proxmox VE supports Proxmox Virtual Environment versions 8.2-9 installed using the official ISO image provided by Proxmox. Veeam Plug-in for Proxmox VE requires at least one file-level storage configured in Proxmox Virtual Environment.
Veeam Software	Veeam Backup & Replication version 13.0.1.180 or later must be deployed on the backup server.
Workers	Workers process backup workload and distribute backup traffic when transferring data to backup repositories. If you deploy a worker using the default configuration, the following compute resources will be allocated: • *CPU*: 6 vCPU • *Memory*: 6 GB RAM • *Disk Space*: 100 GB for product installation and logs With the default configuration, the worker can handle up to 4 concurrent backup and restore tasks. While deploying a new worker or editing settings of an existing one, you can increase the maximum number of concurrent tasks. However, you must allocate 1 vCPU and 1 GB RAM for each additional task. When configuring the maximum number of concurrent tasks, you must also take into account the network traffic throughput in your virtual infrastructure.

IMPORTANT

Workers are backup infrastructure components that are preconfigured for optimal performance. That is why you must not install any software on VMs running as workers or make any configuration changes to them unless you are requested by Veeam Customer Support.

Version Compatibility

The following table lists compatible versions of Veeam Backup & Replication and Veeam Plug-in for Proxmox VE.

Product Release	Veeam Plug-in for Proxmox VE Build	Veeam Backup & Replication Build	Worker OS Version
3	13.3.0.237	13.0.1.180	Rocky Linux 9.2

Product Release	Veeam Plug-in for Proxmox VE Build	Veeam Backup & Replication Build	Worker OS Version
2	13.2.0.457	13.0.0.4967	Rocky Linux 8.10
1.5	12.1.5.17	12.3.2.3617	
1.3	12.1.3.217	12.3.2.3617 12.3.1.1139 12.3.0.310	
1.1	12.1.1.1024	12.3.0.310 12.2.0.334	

Considerations and Limitations

When you plan to use Veeam Plug-in for Proxmox VE, keep in mind the following limitations and considerations.

Configuration

When configuring Veeam Plug-in for Proxmox VE, consider the following:

- Veeam Plug-in for Proxmox VE supports Proxmox Virtual Environment deployments created using the official ISO image provided by Proxmox only.
- Veeam Plug-in for Proxmox VE supports only the `/bin/bash` shell to perform management operations with the Proxmox VE server.
- Veeam Plug-in for Proxmox VE supports custom certificates installed on the Proxmox VE server if either of the following conditions is met:
  - The full certificate chain has been uploaded to the Proxmox VE server.
  - The backup infrastructure components, such as workers, are able to connect to the CA server and validate the certificate.
- Veeam Plug-in for Proxmox VE does not support Online Certificate Status Protocol (OCSP) certificates to access the Proxmox VE server.
- Veeam Plug-in for Proxmox VE does not support credentials of the SSH Private Keys type to access the Proxmox VE server.
- Veeam Plug-in for Proxmox VE does not support user accounts with multi-factor authentication to access the Proxmox VE server.
-
- The Proxmox VE server must be able to establish a direct IP connection to the backup server. Connections through NAT gateways are not supported.
- Before you add a Proxmox VE server to the backup infrastructure, ensure that it has been assigned a unique Proxmox VE system UUID and its name does not contain an FQDN.
- If you want to protect VMs that reside in a Proxmox VE cluster, all nodes of these cluster must be added to the backup infrastructure separately. Adding clusters as standalone entities is not supported.
- After you add nodes of a cluster to the backup infrastructure, you must not change the name of the cluster in the Proxmox VE administration portal.
- After you make changes to your Proxmox VE environment (for example, you migrate a VM between cluster nodes), these changes may not appear in Veeam Backup & Replication immediately — the data synchronization process between the backup server and the Proxmox VE server may take up to 15 minutes to complete. You can speed up the data synchronization process by rescanning the Proxmox VE server.

Backup Repositories

When managing backup repositories, consider that Veeam Plug-in for Proxmox VE does not support storing backups in HPE Cloud Bank Storage repositories. However, you can use them for storing copies of backups created with Veeam Plug-in for Proxmox VE.

Workers

When configuring workers, consider the following:

- The *default* local storage must be enabled on all hosts where worker VMs will be deployed. If you cannot use the default storage in your environment, contact Veeam Customer Support.
- The storage where system files of workers will be stored must support snapshots.
- For VLAN-tagged environments, the VLAN ID must be manually assigned to worker VMs after they are deployed or redeployed.

Backup

When protecting Proxmox VE resources, consider the following:

- Veeam Plug-in for Proxmox VE does not support backup of LXC containers.
- Veeam Plug-in for Proxmox VE does not support backup of VM templates.
- Veeam Plug-in for Proxmox VE does not support backup of VMs created from templates as linked clones. Backup of full clones is supported.
- Veeam Plug-in for Proxmox VE does not support backup of VMs with the same BIOS UUID.
- Veeam Plug-in for Proxmox VE does not support backup of iSCSI disks. If iSCSI disks are attached to a VM included into a backup job, these disks will be skipped from processing.
- Veeam Plug-in for Proxmox VE does not support backup of directly attached (passthrough) disks. If such disks are attached to a VM included into a backup job, these disks will be skipped from processing.
- Veeam Plug-in for Proxmox VE does not support backup of VM permissions granted to users, user groups and API tokens.
- Veeam Plug-in for Proxmox VE does not support backup of VMs that store their disks in the BTRFS and custom storage. All other Proxmox VE storage types are supported.
- The number of concurrent backup operations performed in each storage is limited to 4 to avoid excessive load on the production environment. To change the limit, contact Veeam Customer Support.
- Veeam Plug-in for Proxmox VE does not support VM replication.

Restore

When restoring Proxmox VE resources, consider the following:

- Starting from version 9, Proxmox VE supports the snapshot-as-volume-chain functionality for some storage types. Since this functionality is available for VMs using QEMU version 10 and later, Veeam Plug-in for Proxmox VE is not able to restore VMs using an earlier QEMU version to a storage with the *Allow Snapshots as Volume-Chain* setting enabled. To work around the issue, see this Veeam KB article.
- Veeam Plug-in for Proxmox VE supports only file-level restore from backups stored in HPE StoreOnce Cloud Bank Storage repositories.
- Veeam Plug-in for Proxmox VE supports Instant Recovery of Proxmox VE VMs to the VMware environment with the following limitations:
  - UEFI VMs with MBR cannot be restored.

    - Restored VMs may have an incorrect number of cores per vCPU assigned.
    - Static IP addresses are not restored for Windows VMs.
- Veeam Plug-in for Proxmox VE does not support restore of VMs to the BTRFS and custom storage.

Account Permissions

The accounts used to deploy and administer backup infrastructure components must have the following permissions.

Backup Server Windows Account Permissions

The account used to install Veeam Backup & Replication on a Windows-based machine must have the following permissions.

Account	Required Permission
Setup Account	The account used to install Veeam Backup & Replication and Veeam Plug-in for Proxmox VE must have the Local Administrator permissions on the backup server.
Veeam Backup & Replication User Account	The account used to run Veeam Backup & Replication services must be a *LocalSystem* account or must have the Local Administrator permissions on the backup server.

Proxmox VE Server Permissions

The administrator account that the backup server uses to access the Proxmox VE server must have the *root* or elevated to *root* privileges . For more information on system permissions, see Proxmox VE documentation.

Performing Guest Processing

To allow Veeam Backup & Replication to create application-consistent backups of Windows- and Linux-based VMs, the accounts that will be used to perform guest processing operations (such as transaction log truncation and guest file indexing) must have the permissions listed in this section.

> **NOTE**
>
> The Veeam Backup & Replication console does not provide a possibility to restore application data from application-consistent backups — you can do this using Veeam Explorers only. To see the list of permissions that must be granted to accounts that will be used to perform the restore operations, see the Veeam Explorers User Guide.

Backup Permissions for Windows-Based VMs

For Windows-based VMs, you must choose an account that has administrator privileges. Note that the *Log on as a batch job* permission must be granted to the account and the *Deny log on as a batch job* policy must not be defined. Other permissions depend on applications that you plan to back up:

Application	Required Permission
Microsoft SQL Server	To back up Microsoft SQL Server data, the user whose account you plan to use must have the following permissions: • SQL Server instance-level role: *public* and *dbcreator*. • Database-level roles and roles for the model system database: *db_backupoperator*, *db_denydatareader*, *public*; for the master system database — *db_backupoperator*, *db_datareader*, *public*; for the msdb system database — *db_backupoperator*, *db_datareader*, *public, db_datawriter*. • Securables: *view any definition*, *view server state, connect SQL*. **Tip**: If you do not want to assign the permissions gradually, use an account that has local Administrator permissions on the target VM and system Administrator permissions (with the Sysadmin role) on the target Microsoft SQL Server.
Microsoft Active Directory	The account used to back up Microsoft Active Directory data or a Domain Controller server must be a member of the built-in *Administrators* group. The account used to back up a Read-Only Domain controller can have permissions of a delegated RODC administrator account. For more information, see Microsoft Docs.
Microsoft Exchange	The account used to back up Microsoft Exchange data must have the local Administrator permissions on the machine where Microsoft Exchange is installed.
Oracle	The account used to communicate with VM guest OSes must be a member of both the *Local Administrator* group and the *ORA_DBA* group (if OS authentication is used). In addition, if *ASM* is used, then such an account must be a member of the *ORA_ASMADMIN* group (for Oracle 12 and higher). The account used to back up Oracle databases must have have the following permissions: • Oracle account with SYSDBA privileges. You can use, for example, the SYS Oracle account or any other Oracle account that has been granted SYSDBA privileges. • Account specified for guest processing. That is, the **Use guest credentials** option selected. In this case, the account that was specified at the **Guest Processing** step must be a member of the *ORA_DBA* group.

Application	Required Permission
Microsoft SharePoint	The account used to back up Microsoft SharePoint server data must have the Farm Administrator role. The account used to back up Microsoft SQL databases of the Microsoft SharePoint Server must have the same privileges as that of Microsoft SQL Server.

TIP

The account must be specified either in the *DOMAIN|USERNAME* (for Active Directory accounts) or in the *HOST|USERNAME* (for local user accounts) format.

Backup Permissions for Linux-Based VMs

For Linux-based VMs, you must choose an account of a root user or a user elevated to root. Note that the account must have the `/home` directory created. Other permissions depend on applications that you plan to back up:

Application	Required Permission
Oracle	The account used to back up Oracle databases must have have the following permissions: • Oracle account with SYSDBA privileges. You can use, for example, the SYS Oracle account or any other Oracle account that has been granted SYSDBA privileges. • Account specified for guest processing. That is, the **Use guest credentials** option selected. In this case, the account that was specified at the **Guest Processing** step must be a member of the *OSASM*, *OSDBA* and *OINSTALL* groups. **Note**: To perform guest processing of Oracle databases running on Linux servers, make sure that the `/tmp` directory is mounted with the `exec` option. Otherwise, you will get a permission denial error.
PostgreSQL	The account used to back up PostgreSQL instances must have superuser privileges for the PostgreSQL instance. For more information, see PostgreSQL documentation. The following permissions must be granted to access the folder used as a temporary location for archive logs: • The user running the PostgreSQL instance must have *read*, *write*, and *execute* (*rwx*) permissions. • The user selected in the backup job settings to access the guest OS must have *read* and *execute* (*rx*) permissions.

Ports

Veeam Backup & Replication automatically creates firewall rules for the ports required to allow communication between the Proxmox VE server, workers and the backup server.

Workers

The following table describes network ports that must be open to ensure proper communication of workers with other backup infrastructure components.

From	To	Protocol	Port	Notes
Worker	Proxmox VE server	TCP/HTTPS	8006	Used to communicate with the REST API service running on the Proxmox VE server.
	Proxmox VE server	SSH	22	Used to communicate with Proxmox VE server.
	Backup server	TCP	10006	Used to communicate with the backup server.
	Veeam backup repository or gateway server	TCP	2500-3300	Default range of ports used as transmission channels for jobs and restore sessions. For each TCP connection that a job uses, one port from this range is assigned.
	Veeam Update Repository (repository.veeam.com) Amazon CloudFront (cloudfront.net, amazonaws.com)	TCP/HTTPS	443	Used to download worker deployment packages. **Note**: Veeam Update Repository uses the Amazon CloudFront service to distribute traffic when downloading product updates.

Backup Server

The following table describes network ports that must be open to ensure proper communication of the backup server with other backup infrastructure components.

From	To	Protocol	Port	Notes
	Worker	TCP	19000	Used to communicate with workers.

From	To	Protocol	Port	Notes
Backup server	Worker	TCP/HTTPS	443	Used by the Platform Service to enable communication with the Veeam Updater service on the worker.
	Backup server	TCP/HTTPS	6172	Used by the Platform Service to enable communication with the Veeam Backup & Replication database.
	Proxmox VE server	TCP/HTTPS	8006	Used to communicate with the REST API service running on the Proxmox VE server.
	Proxmox VE server	SSH	22	Used to communicate with the Proxmox VE server.
	FLR helper appliance	TCP	22	Used to connect to the helper appliance during file-level restore.

NOTE

For the list of ports used by the backup server to communicate with backup repositories, see the Veeam Backup & Replication User Guide, section Used Ports.

Guest Processing Components

Connections with Non-Persistent Runtime Components

The following tables describe network ports that must be opened to ensure proper communication of the backup server and backup infrastructure components with the non-persistent runtime components deployed inside the VM guest OS for application-aware processing and indexing.

From	To	Protocol	Port	Notes
Backup server	VM guest OS (Linux)	TCP	22	Default SSH port used as a control channel.
		TCP	2500 to 3300	Default range of ports used as transmission channels for log shipping.
		TCP	6190	Used for communication with the guest interaction proxy.

From	To	Protocol	Port	Notes
	Guest interaction proxy	TCP	6290	Used as a control channel for communication with the guest interaction proxy.
		TCP	445	Port used as a transmission channel.
Guest interaction proxy	VM guest OS (Microsoft Windows)	TCP	445 135	Required to deploy the runtime coordination process on the VM guest OS.
		TCP	2500 to 3300	Default range of ports used as transmission channels for log shipping.
		TCP	49152 to 65535	Dynamic RPC port range for Microsoft Windows 2008 and later. For more information, see this Microsoft KB article. Used by the runtime process deployed inside the VM for guest OS interaction. **Note:** If you use default Microsoft Windows firewall settings, you do not need to configure dynamic RPC ports. During setup, Veeam Backup & Replication automatically creates a firewall rule for the runtime process. If you use firewall settings other than default ones or application-aware processing fails with the *"RPC function call failed"* error, you need to configure dynamic RPC ports. For more information on how to configure RPC dynamic port allocation to work with firewalls, see this Microsoft KB article.
VM guest OS	Guest interaction proxy or backup server	TCP	2500 to 3300	Default range of ports used as transmission channels for log shipping.

Log Shipping Components

The following tables describe network ports that must be opened to ensure proper communication between log shipping components.

- Log Shipping Server Connections
- MS SQL Guest OS Connections

- Oracle Guest OS Connections
- PostgreSQL Guest OS Connections

Log Shipping Server Connections

From	To	Protocol	Port	Notes
Backup server	Log shipping server	TCP	445 135	Required for deploying Veeam Backup & Replication components.
		TCP	6160	Default port used by Veeam Installer Service.
		TCP	6162	Default port used by Veeam Data Mover Service.
		TCP	49152 to 65535	Dynamic RPC port range for Microsoft Windows 2008 and later. For more information, see this Microsoft KB article. **Note:** If you use default Microsoft Windows firewall settings, you do not need to configure dynamic RPC ports. During setup, Veeam Backup & Replication automatically creates a firewall rule for the runtime process. If you use firewall settings other than default ones or application-aware processing fails with the *"RPC function call failed"* error, you need to configure dynamic RPC ports. For more information on how to configure RPC dynamic port allocation to work with firewalls, see this Microsoft KB article.
Log shipping server	Backup repository	TCP	2500 to 3300	Default range of ports used for communication with a backup repository and transfer log backups.

MS SQL Guest OS Connections

From	To	Protocol	Port	Notes
Guest interaction proxy	MS SQL VM guest OS	TCP	445 135	Required for deploying Veeam Backup & Replication components including Veeam Log Shipper runtime component.

From	To	Protocol	Port	Notes
		TCP	2500 to 3300	Default range of ports used for communication with a guest OS.
		TCP	49152 to 65535	Dynamic RPC port range for Microsoft Windows 2008 and later. For more information, see this Microsoft KB article. **Note:** If you use default Microsoft Windows firewall settings, you do not need to configure dynamic RPC ports. During setup, Veeam Backup & Replication automatically creates a firewall rule for the runtime process. If you use firewall settings other than default ones or application-aware processing fails with the *"RPC function call failed"* error, you need to configure dynamic RPC ports. For more information on how to configure RPC dynamic port allocation to work with firewalls, see this Microsoft KB article.
		TCP	6167	Used by the Veeam Log Shipping Service for preparing the database and taking logs.
MS SQL VM guest OS	Guest interaction proxy	TCP	2500 to 3300	Default range of ports used for communication with a guest interaction proxy.
MS SQL VM guest OS	Backup repository	TCP	2500 to 3300	Default range of ports used for communication with a backup repository and transfer log backups. Should be opened if log shipping servers are not used in the infrastructure and the MS SQL server has a direct connection to the backup repository.
MS SQL VM guest OS	Log shipping server	TCP	2500 to 3300	Default range of ports used for communication with a log shipping server and transfer log backups.

Oracle Guest OS Connections

From	To	Protocol	Port	Notes
Guest interaction proxy	Oracle VM guest OS (Microsoft Windows)	TCP	445 135	Required for deploying Veeam Backup & Replication components including Veeam Log Shipper runtime component.
		TCP	2500 to 3300	Default range of ports used for communication with a guest OS.
		TCP	49152 to 65535	Dynamic RPC port range for Microsoft Windows 2008 and later. For more information, see this Microsoft KB article. **Note:** If you use default Microsoft Windows firewall settings, you do not need to configure dynamic RPC ports. During setup, Veeam Backup & Replication automatically creates a firewall rule for the runtime process. If you use firewall settings other than default ones or application-aware processing fails with the *"RPC function call failed"* error, you need to configure dynamic RPC ports. For more information on how to configure RPC dynamic port allocation to work with firewalls, see this Microsoft KB article.
		TCP	6167	Used by the Veeam Log Shipping Service for preparing the database and taking logs.
Backup server	Oracle VM guest OS (Linux)	TCP	22	Default SSH port used as a control channel.
		TCP	2500 to 3300	Default range of ports used for communication with a guest OS.
Oracle VM guest OS	Guest interaction proxy or backup server	TCP	2500 to 3300	Default range of ports used for communication with a guest interaction proxy.

From	To	Protocol	Port	Notes
Oracle VM guest OS	Backup repository	TCP	2500 to 3300	Default range of ports used for communication with a backup repository and transfer log backups. Should be opened if log shipping servers are not used in the infrastructure and the Oracle server has a direct connection to the backup repository.
Oracle VM guest OS	Log shipping server	TCP	2500 to 3300	Default range of ports used for communication with a log shipping server and transfer log backups.

PostgreSQL Guest OS Connections

From	To	Protocol	Port	Notes
Backup server	PostgreSQL VM guest OS	TCP	22	Default SSH port used as a control channel.
		TCP	2500 to 3300	Default range of ports used for communication with a guest OS.
PostgreSQL VM guest OS	Backup server	TCP	2500 to 3300	Default range of ports used for communication with a guest interaction proxy.
PostgreSQL VM guest OS	Backup repository	TCP	2500 to 3300	Default range of ports used for communication with a backup repository and transfer log backups. Should be opened if log shipping servers are not used in the infrastructure and the PostgreSQL server has a direct connection to the backup repository.
PostgreSQL VM guest OS	Log shipping server	TCP	2500 to 3300	Default range of ports used for communication with a log shipping server and transfer log backups.

Licensing

Veeam Plug-in for Proxmox VE is licensed by the number of protected Proxmox VE VMs. Each protected Proxmox VE VM consumes one Veeam Universal License instance from the license scope. A Proxmox VE VM is considered protected if it has a restore point created during the past 31 days.

By default, Veeam Plug-in for Proxmox VE automatically revokes a license instance from a protected VM if no new restore points have been created during the past 31 days. However, you can manually revoke license instances from protected VMs as described in the Veeam Backup & Replication User Guide, section Revoking License.

Obtaining New License

You can obtain the following types of licenses for Veeam Plug-in for Proxmox VE:

- **Evaluation license** is a free license that can be used for product evaluation. The license is valid for 30 days from the moment of the product download.

  To obtain this license, request a trial key on the Veeam downloads page as described in the Veeam Backup & Replication User Guide, section Obtaining and Renewing License.

- **Subscription license** is a paid license with a limited subscription term. The expiration date of the Subscription license is set to the end of the subscription term. The Subscription license term is normally 1-5 years from the license issue date.

  To obtain this license, choose the required subscription term on the Veeam Backup & Replication Pricing page and contact the Veeam Sales Team.

- **Perpetual license** is a paid license without an expiration date. The Perpetual license typically includes one year period of basic support and maintenance that can be extended.

  To obtain this license, contact a reseller in your region.

After you obtain a license, install it on the backup server as described in the Veeam Backup & Replication User Guide, section Installing License.

Using Existing License

If you already use Veeam Backup & Replication and you have spare Veeam Universal License instances on your backup server, they can be used to protect Proxmox VE VMs. You can check the number of available license instances in the Veeam Backup & Replication console as described in the Veeam Backup & Replication User Guide, section Viewing License Information.

Deployment

Starting from version 12.2, the Veeam Backup & Replication solution allows you to add Proxmox VE servers to the backup infrastructure, and to manage data protection and recovery operations for Proxmox VE workloads from a single console.

To access the Veeam Plug-in for Proxmox VE functionality, you can either deploy a new backup server as described in the Veeam Backup & Replication User Guide or use a backup server that already exists in your backup infrastructure if it meets the Veeam Plug-in for Proxmox VE system requirements.

Upgrading to Veeam Plug-In for Proxmox VE 3

You can upgrade Veeam Plug-in for Proxmox VE from version 1.3, 1.5 or 2 to version 3.

Before you start the upgrade process, do the following:

- [Apples only to Veeam Plug-in for Proxmox VE prior to 2] Upgrade your Veeam Backup & Replication server to version 12.3 (12.3.1 or later) as described in the Veeam Backup & Replication User Guide, section Upgrading to Veeam Backup & Replication 12.

  Ensure that Veeam Plug-in for Proxmox VE has upgraded to version 1.3 or 1.5. For more information see the Veeam Plug-in for Proxmox VE User Guide, section Installing Proxmox VE Plug-in Manually.

- Download Veeam Backup & Replication version 13.0.1 from the Veeam downloads page.

To upgrade Veeam Plug-in for Proxmox VE to version 3, do the following:

1. Upgrade your Veeam Backup & Replication server to version 13.0.1 as described in the Veeam Backup & Replication User Guide, section Upgrading to Veeam Backup & Replication 13.
2. Complete the **Components Update** wizard as described in the Veeam Backup & Replication User Guide, section Server Components Upgrade.

Installing Veeam Plug-In for Proxmox VE Manually

[This section applies only to Windows-based backup servers]

The pre-installed plug-in that comes with the default installation package of Veeam Backup & Replication allows you to protect Proxmox VE resources. However, you may require to install a new plug-in version on the backup server manually if some updates and patches become available.

> **NOTE**
>
> If you use a remote Veeam Backup & Replication console, you do not need to install Veeam Plug-in for Proxmox VE on the workstation where the remote Veeam Backup & Replication console is deployed.

To install Veeam Plug-in for Proxmox VE, do the following:

1. Log in to the backup server using an account with the local Administrator permissions.
2. Download a product installation file from your Veeam download page.
3. Open the downloaded archive file and launch the installation file.

   Before proceeding with installation, the installer will check whether you have Microsoft .NET Core Runtime installed on the backup server. In case the required version is missing, the installer will offer to install it automatically. To do that, click **OK**.

4. At the **License Agreement** step of the **Veeam Plug-In Plug-in for Proxmox Virtual Environment** setup wizard, read and accept both the Veeam license agreement, licensing policy, the 3rd party components and required software license agreement. If you reject the agreements, you will not be able to continue installation.

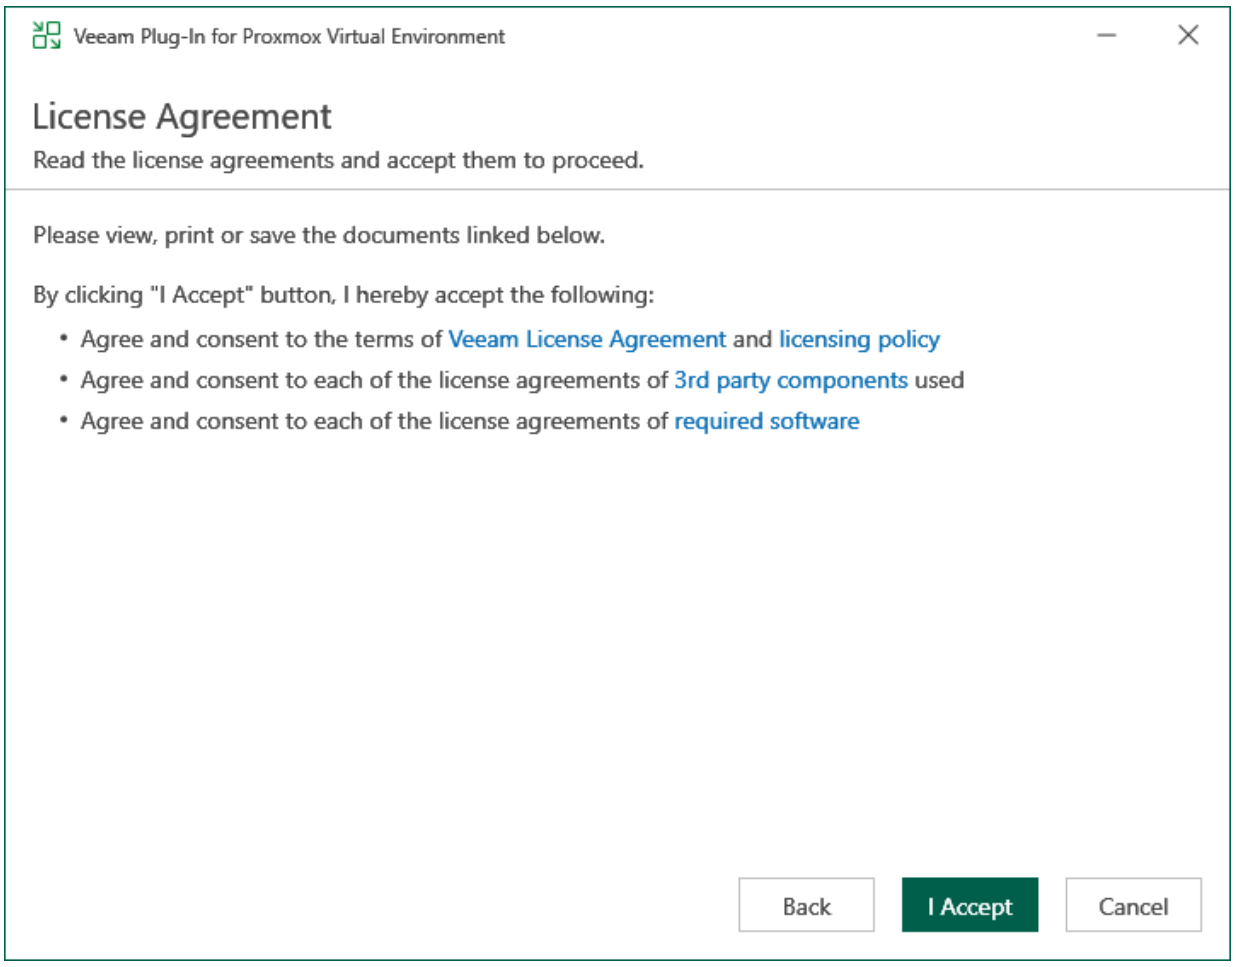

5. At the **Data Location** step of the wizard, you can change the installation directory if necessary.

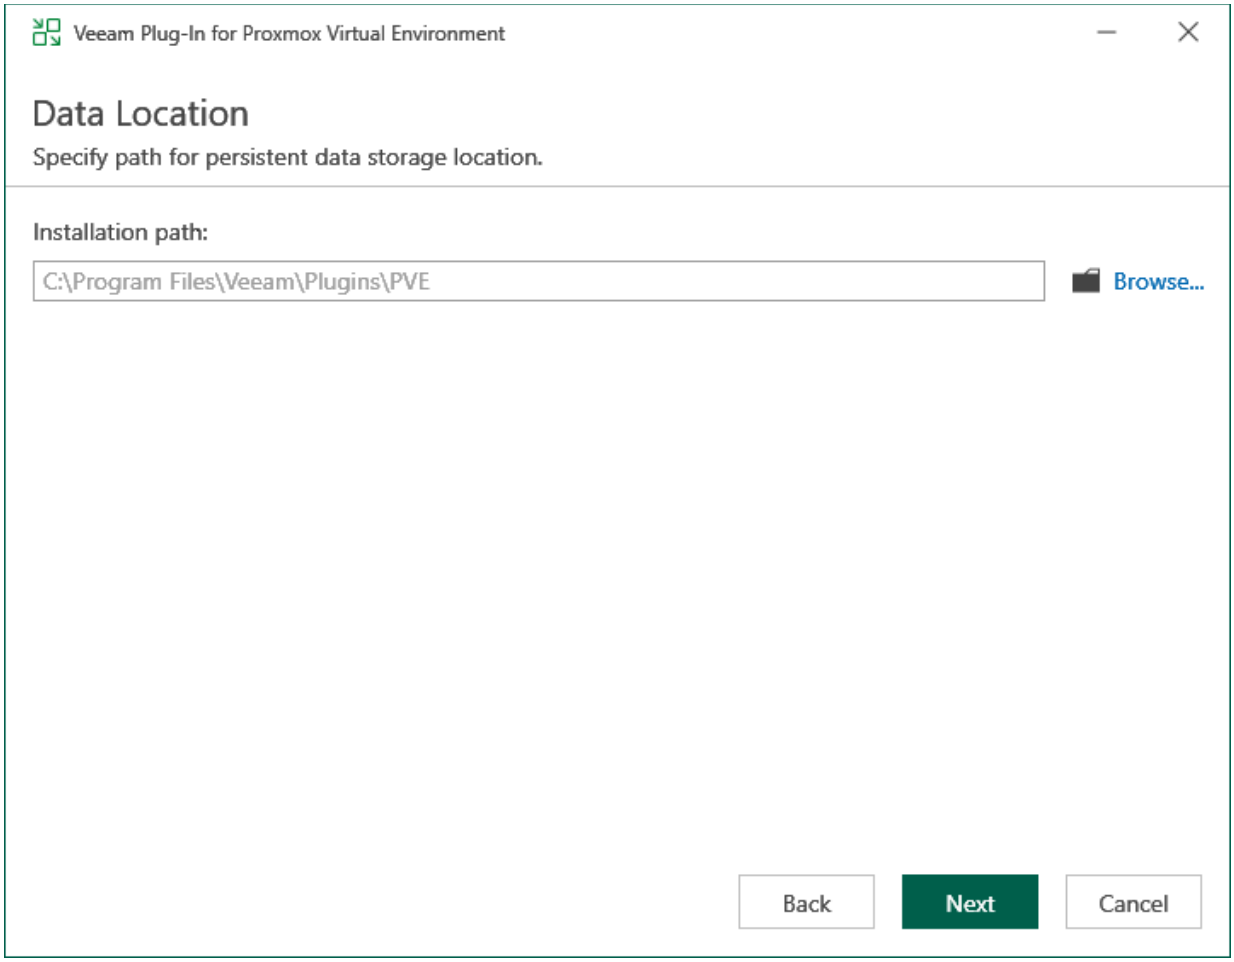

6. At the **Ready to Install** step of the wizard, click **Install** to begin installation.

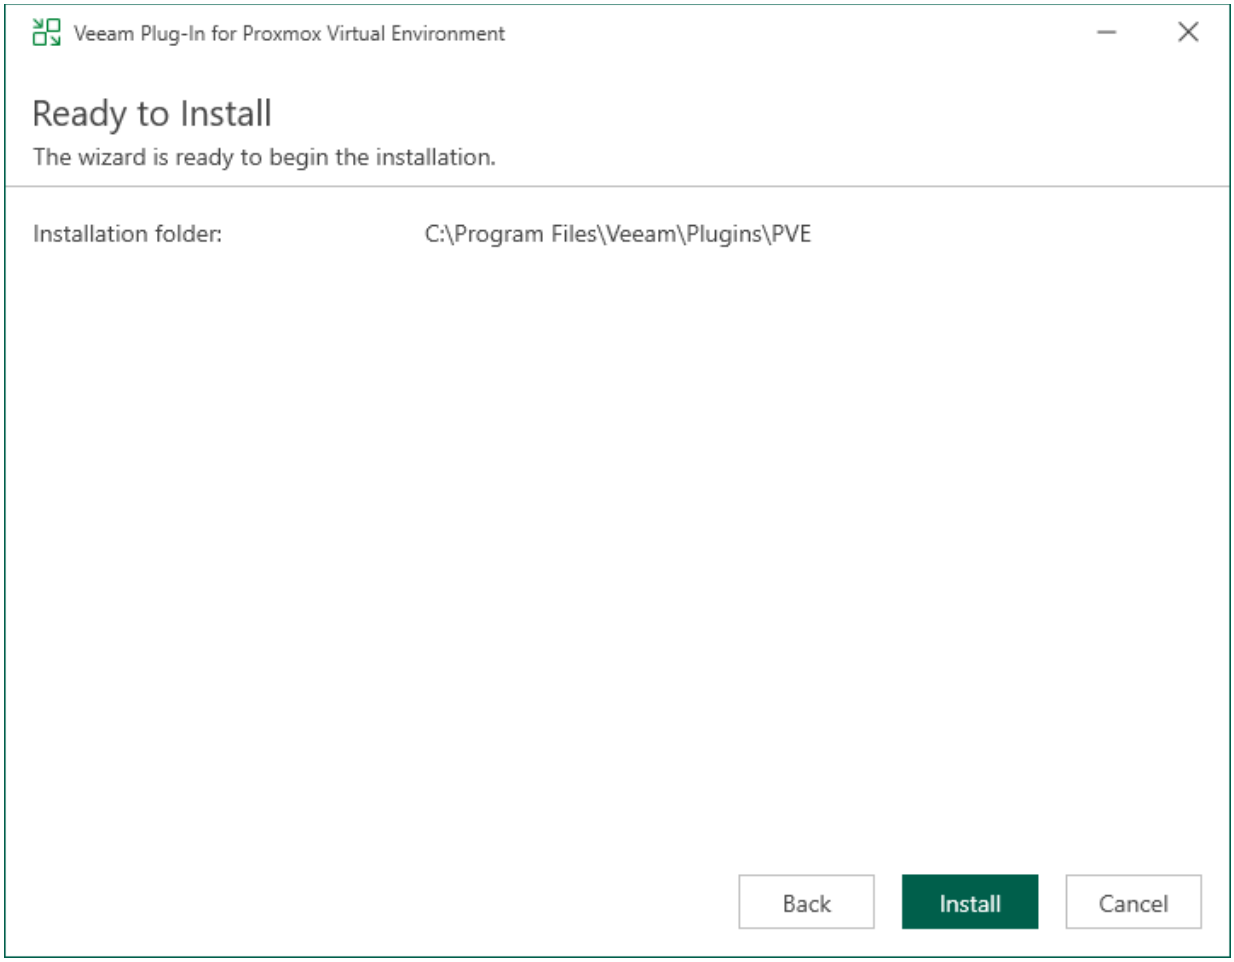

Uninstalling Veeam Plug-In for Proxmox VE Manually

Before you uninstall Veeam Plug-in for Proxmox VE, it is recommended to remove all configured workers from the backup infrastructure.

To uninstall Veeam Plug-in for Proxmox VE, do the following:

1. Log in to the backup server using an account with the Local Administrator permissions.
2. Open the **Start** menu and click the **Settings** icon.
3. In the **Settings** window, navigate to **System** > **Apps and Features**.
4. In the program list, select **Veeam Plug-in for Proxmox Virtual Environment**. Then, click **Uninstall**.
5. In the opened window, click **Remove**.

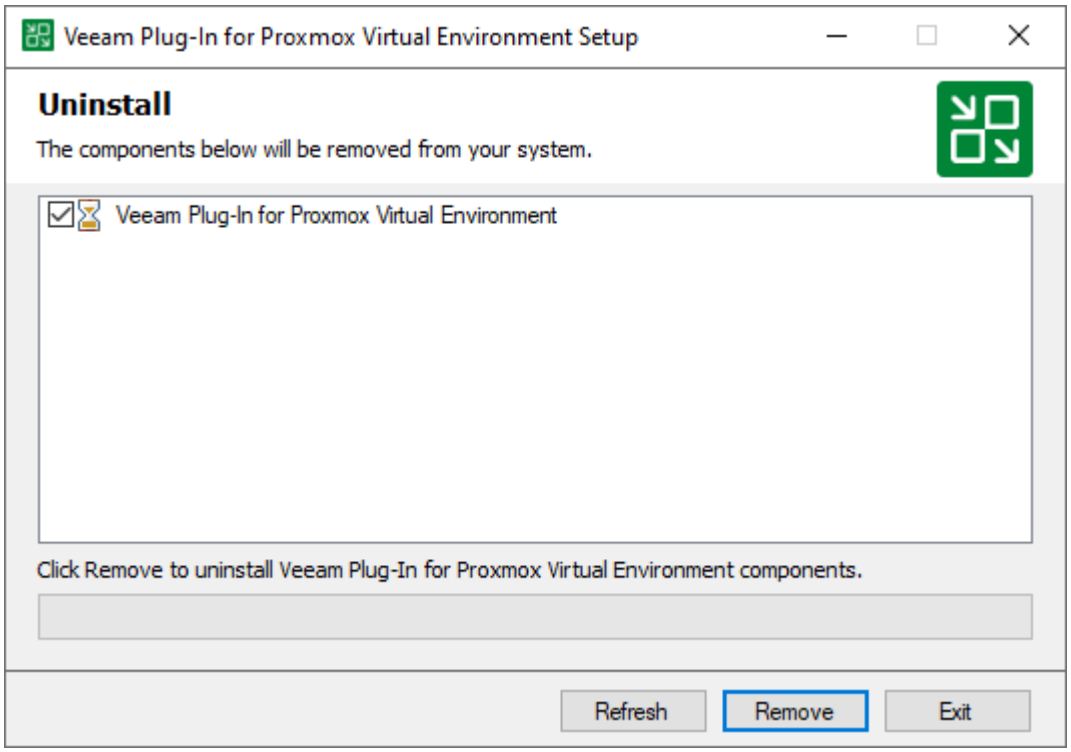

Configuring Veeam Plug-In for Proxmox VE

To start working with Veeam Plug-in for Proxmox VE, perform a number of steps for its configuration:

1. Configure backup repositories where Veeam Backup & Replication will store backups of Proxmox VE VMs.
2. Connect to Proxmox VE servers that administer Proxmox VE resources you want to protect.
3. Deploy workers that will transfer backup traffic.
4. [Optional] Configure email settings and notifications.

Configuring Backup Repositories

A backup repository is a storage location where Veeam Backup & Replication keeps backup files. By default, the backup server performs the role of a backup repository. To keep your backups in another storage location, you can configure the following types of repositories:

- **Direct attached storage**: Microsoft Windows and Linux virtual and physical machines. Hardened repositories based on Linux servers are also supported.
- **Network attached storage**: CIFS (SMB) shares and NFS shares.
- **Deduplicating storage appliances**: ExaGrid, Quantum DXi, Dell Data Domain, HPE StoreOnce, Fujitsu ETERNUS, Infinidat InfiniGuard.
- **Cloud object storage**: Amazon S3, S3 compatible, Google Cloud, Wasabi Cloud Storage, Veeam Data Cloud Vault, IBM Cloud and Microsoft Azure Blob.

To combine repositories of different types in one repository, you can configure a scale-out backup repository and add any of supported repositories to its performance tier.

For Linux server, Microsoft Windows server, SMB share, ExaGrid, Quantum DXi, Fujitsu ETERNUS and Infinidat InfiniGuard repositories, you can enable the Fast Clone technology that increases the speed of synthetic backup creation and transformation, reduces disk space requirements and decreases the load on storage devices. With this technology, Veeam Backup & Replication references existing data blocks on volumes instead of copying data blocks between files. Data blocks are copied only when files are modified. To learn how to configure a repository to enable this functionality, see the Veeam Backup & Replication User Guide, section Fast Clone.

> **IMPORTANT**
>
> Veeam Plug-in for Proxmox VE does not support storing backups in HPE Cloud Bank Storage repositories. However, you can use them for storing copies of backups created with Veeam Plug-in for Proxmox VE.

Connecting Proxmox VE Server

A Proxmox VE server is a Proxmox VE standalone host or cluster that allows the backup server to access Proxmox VE resources such as VMs, storage and networks. After you add a Proxmox VE server to the backup infrastructure, you will be able to deploy workers and to manage data protection tasks for Proxmox VE VMs.

> **IMPORTANT**
>
> If you want to add a Proxmox VE cluster to the backup infrastructure, consider the following:
>
> - Each node of the cluster must be added to the backup infrastructure separately.
> - The name of the cluster must not be changed in the Proxmox VE environment after you add it to the backup infrastructure.
> - The type of the Proxmox VE server (standalone host or cluster) must not be changed in the Proxmox VE environment after you add it to the backup infrastructure.
>
>   If you want to add the host that was already connected to the backup server as a node to a Proxmox VE cluster, first remove the host from the backup infrastructure and then add it again. Keep in mind that Veeam Backup & Replication will treat it as a new sever and will start new backup chains for VMs.

Adding Proxmox VE Server to Backup Infrastructure

To add a Proxmox VE server to the backup infrastructure, do the following:

1. Launch the New Proxmox VE Server wizard.
2. Specify the Proxmox VE server domain name or IP address.
3. Enter credentials to access the Proxmox VE server.
4. Select storage for snapshots.
5. Apply Proxmox VE server settings.
6. Finish working with the wizard.

Step 1. Launch New Proxmox VE Server Wizard

To launch the **New Proxmox VE Server** wizard, do the following:

1. In the Veeam Backup & Replication console, open the **Backup Infrastructure** view.
2. In the inventory pane, select **Managed Servers**.
3. On the ribbon, click **Add Server**.
4. In the **Add Server** window, select **Virtualization Platforms**.
5. In the **Virtualization Platforms** window, select **Proxmox VE** to launch the **New Proxmox VE Server** wizard.

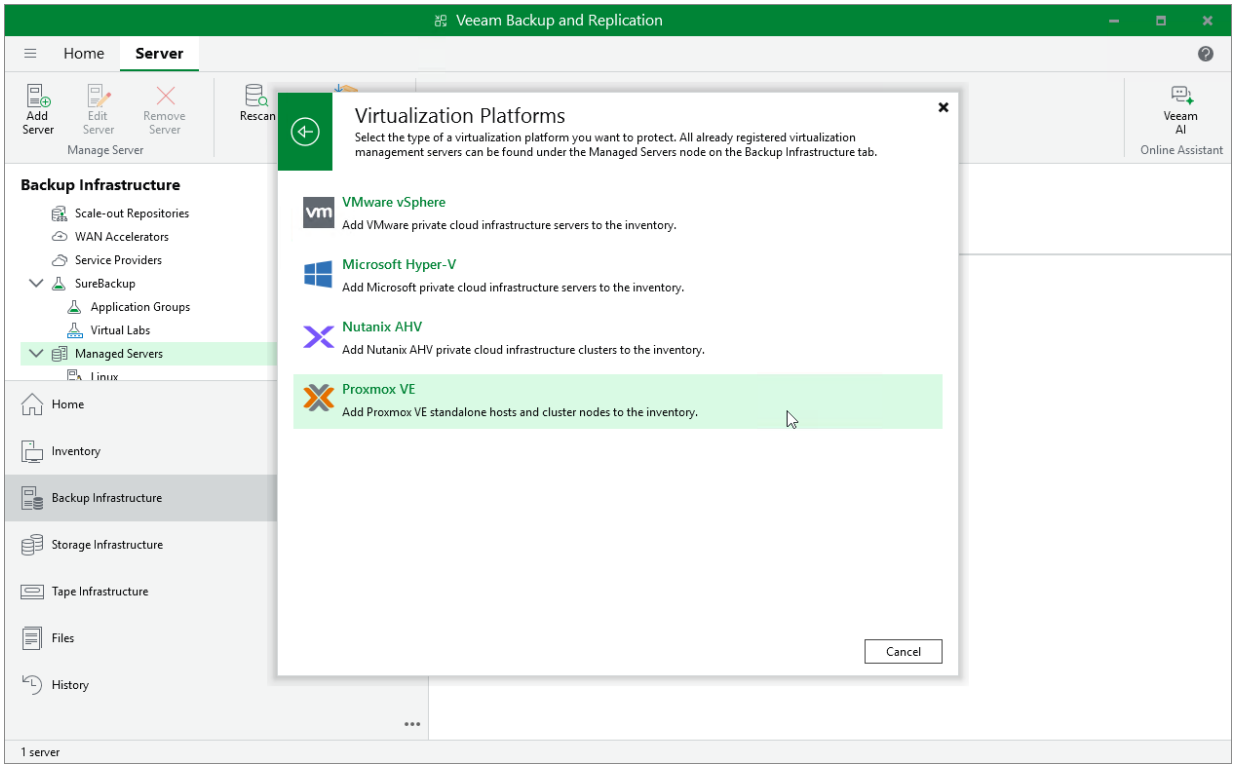

Step 2. Specify Domain Name or IP Address of Proxmox VE server

At the **Name** step of the wizard, do the following:

1. In the **DNS name or IP address** field, enter the FQDN or IP address of the Proxmox VE standalone host or cluster node.
2. In the **Description** field, provide a description for future reference. The field already contains a default description with information about the user who added the manager, date and time when the manager was added.

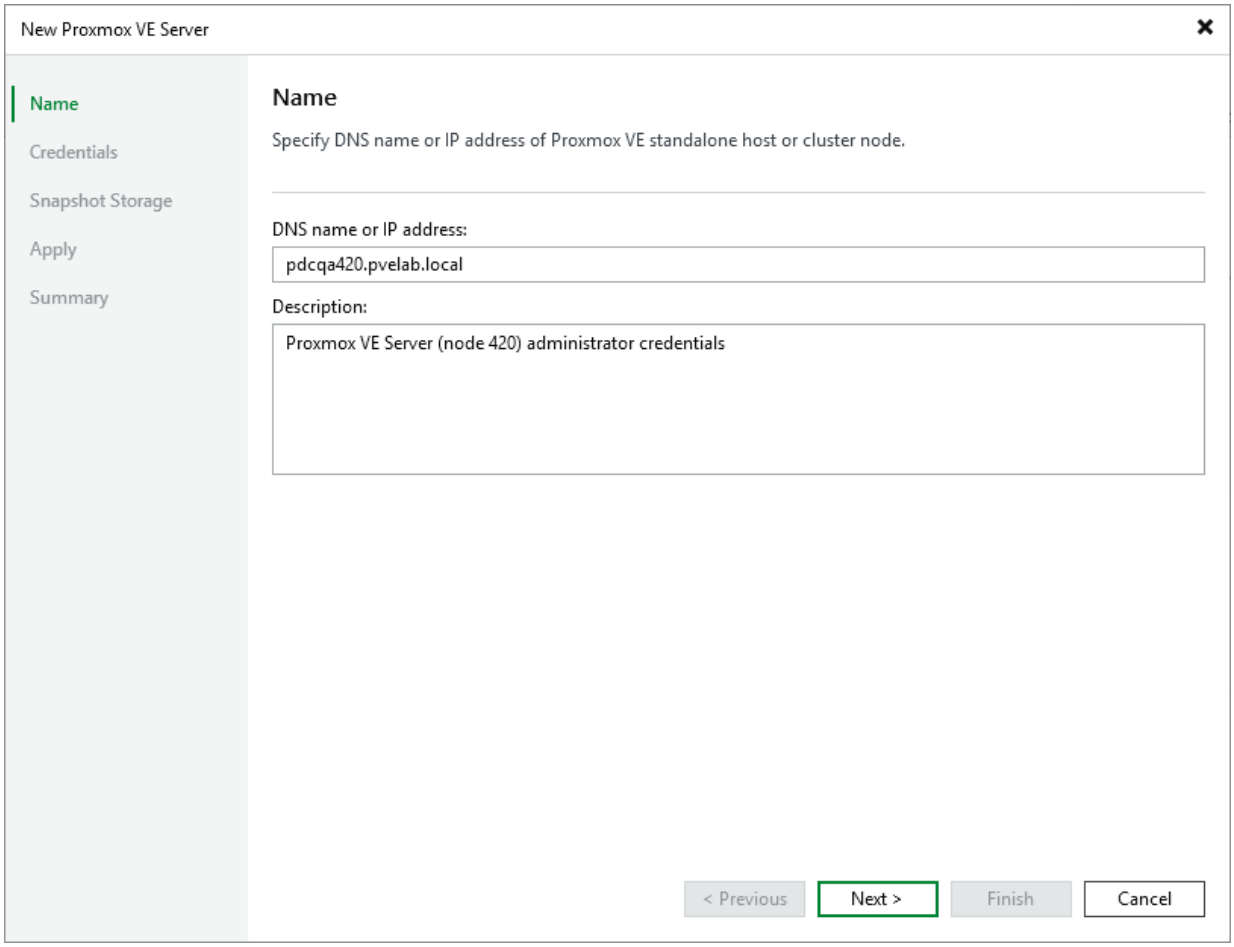

Step 3. Enter Credentials

> **NOTE**
>
> In versions prior to 1.3, Veeam Plug-in for Proxmox VE used credentials of the Standard type to connect to Proxmox VE servers. In version 1.3, those credentials are automatically converted to credentials of the SSH type. However, keep in mind that credentials of the SSH Private Keys type are not supported.

At the **Credentials** step of the wizard, specify credentials of an account that will be used to access the Proxmox VE server — it can be either an account of a root user or an account of a user elevated to root (the latter option is recommended for security reasons).

For credentials to be displayed in the **Credentials** list, they must be added to the Credentials Manager as described in the Veeam Backup & Replication User Guide, section SSH Credentials. If you have not added the necessary credentials to the Credentials Manager beforehand, you can do this without closing the **New Proxmox VE Server** wizard. To do that, click either the **Manage accounts** link or the **Add** button, and specify the user name, password and description in the **Credentials** window.

> **TIP**
>
> If you want to use an account of a user elevated to root, you must also select the **Elevate account privileges automatically** check box. Consider that Veeam Plug-in for Proxmox VE ignores the **Add account to the sudoers file** and **Use "su" if "sudo" fails** options for security reasons.

After you click **Next**, the backup server will connect to the Proxmox VE server and check its TLS certificate. If the certificate is not installed on the backup server, the **Certificate Security Alert Window** will display a warning notifying that secure communication cannot be guaranteed. To allow the backup server to connect to the Proxmox VE server using the certificate, click **Continue**.

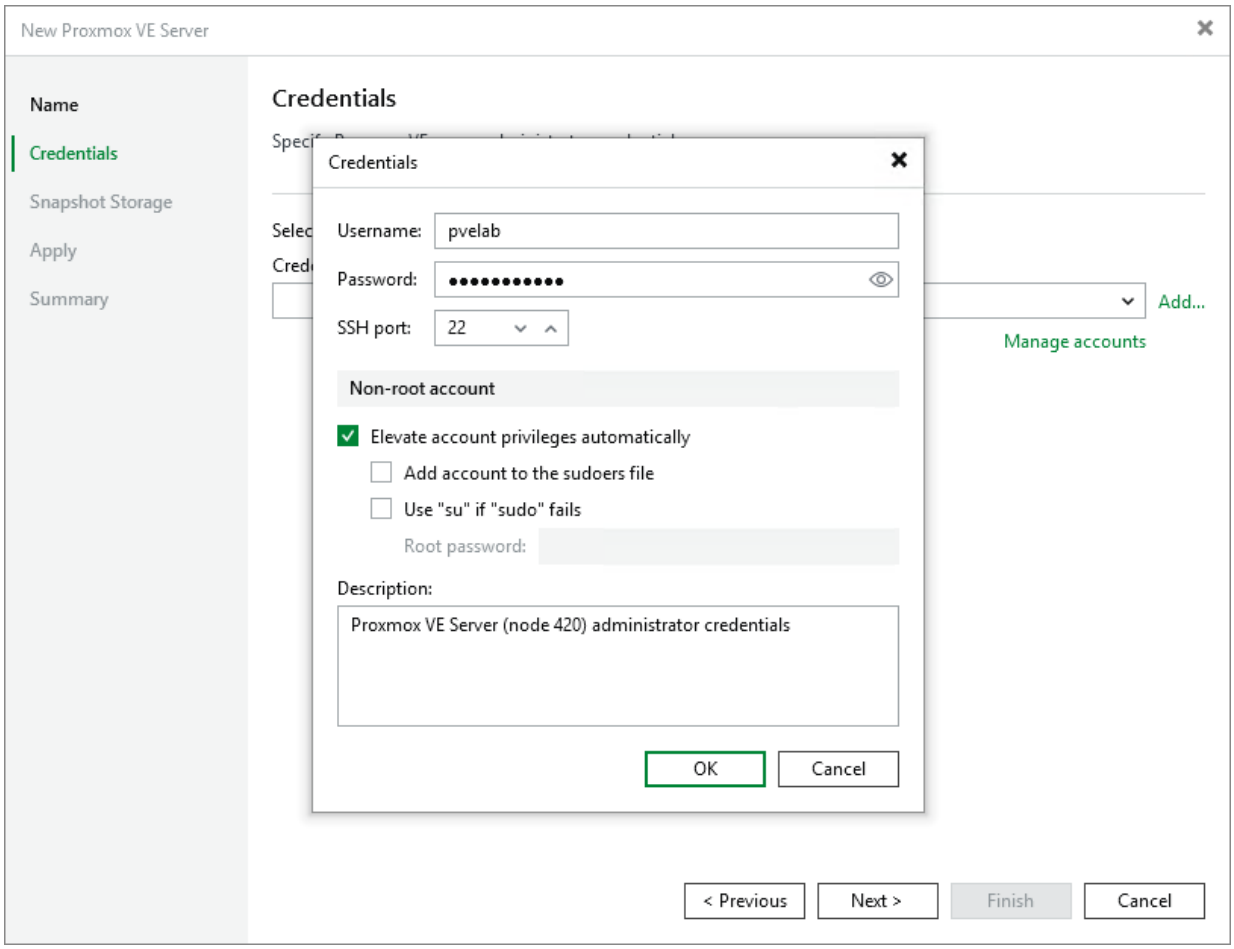

Step 4. Configure Storage Settings

At the **Snapshot Storage** step of the wizard, choose whether you want to keep snapshots of processed VMs in specific storage or in the largest file-level storage available on the connected Proxmox VE server — but only in case the original VM storage does not support snapshots.

For more information on storage that supports snapshots, see Proxmox VE documentation.

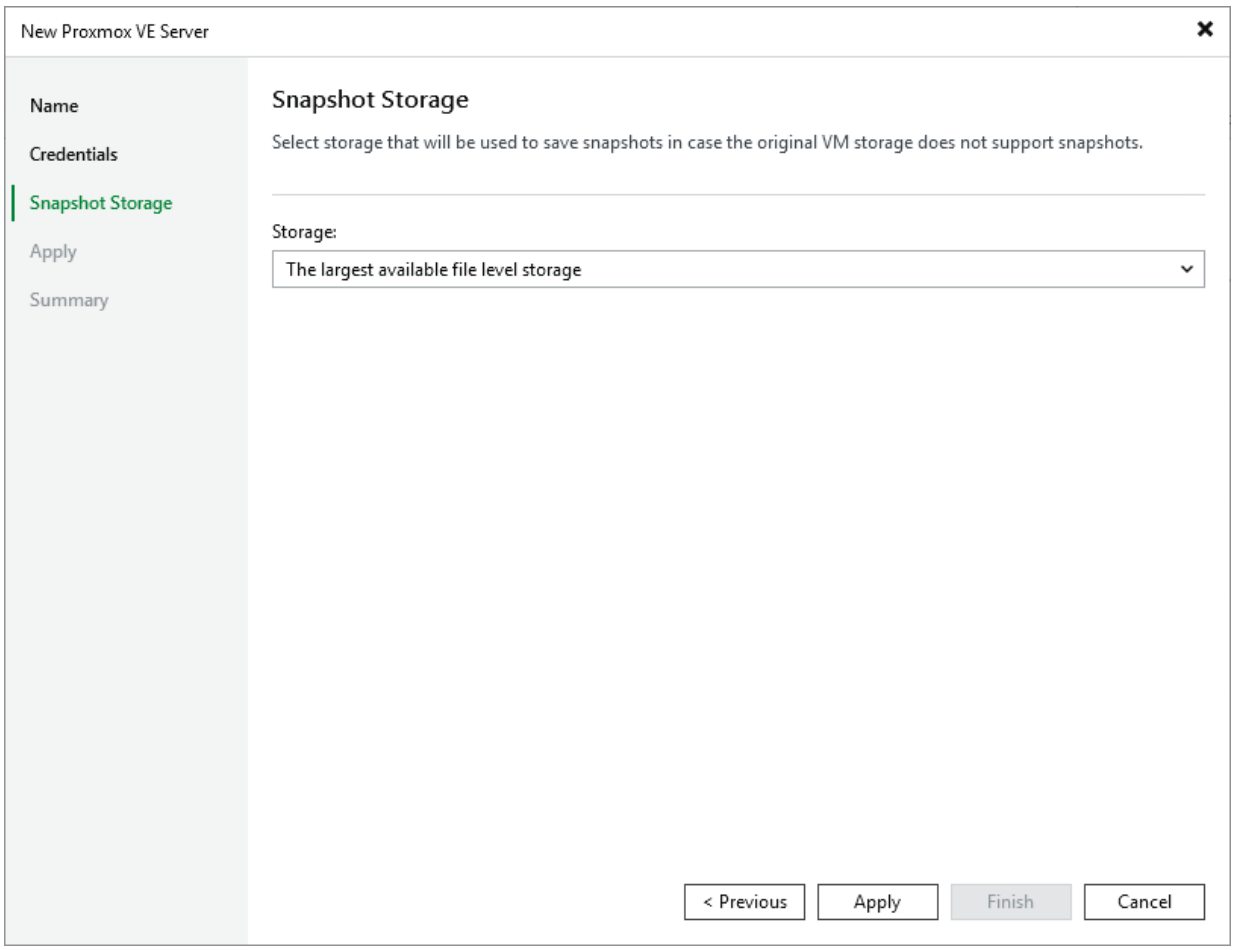

Step 5. Apply Settings

At the **Apply** step of the wizard, wait until the Proxmox VE server is added to the backup infrastructure and then click **Next**.

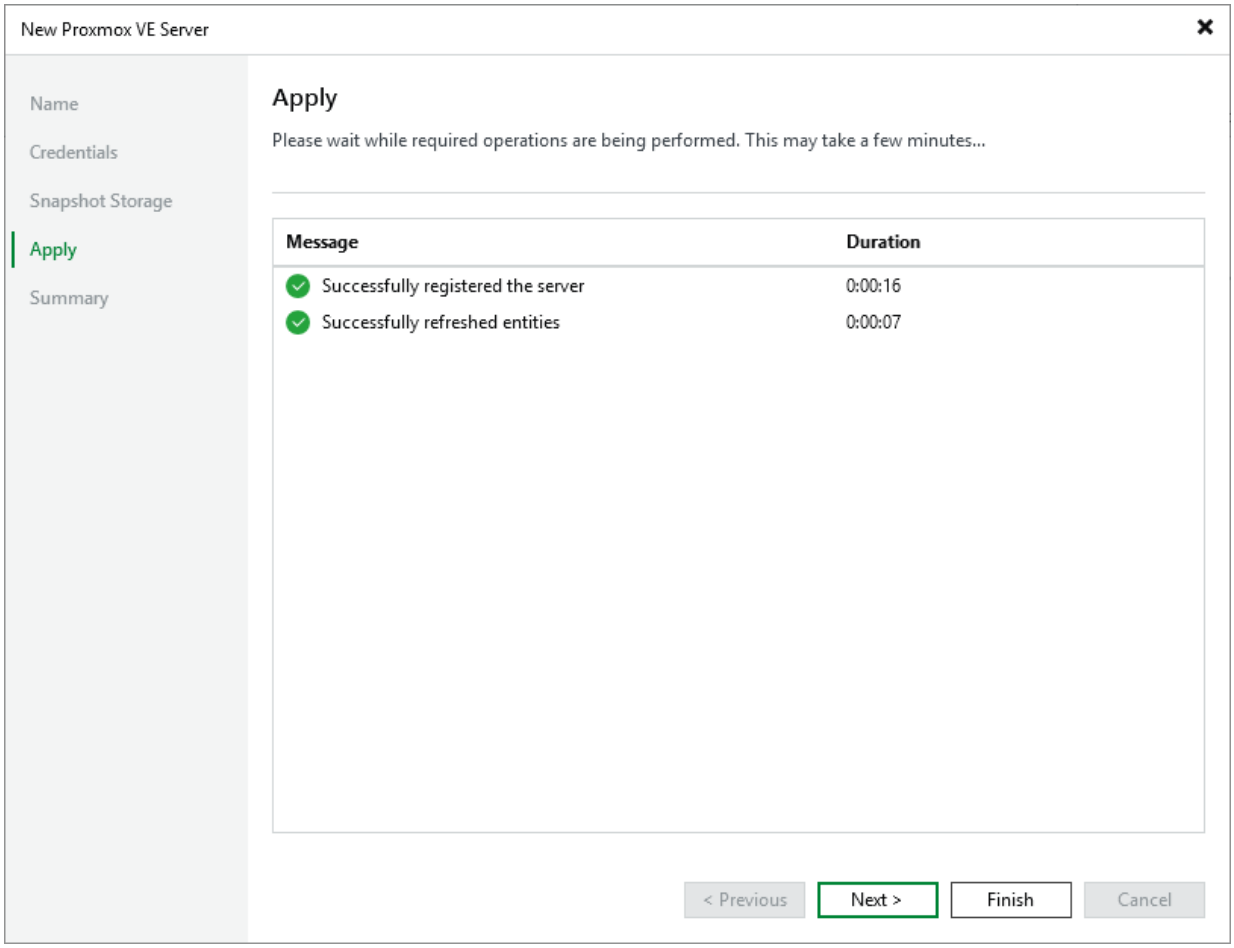

Step 6. Finish Working with Wizard

At the **Summary** step of the wizard, check that the Proxmox VE server has been successfully added and click **Finish**.

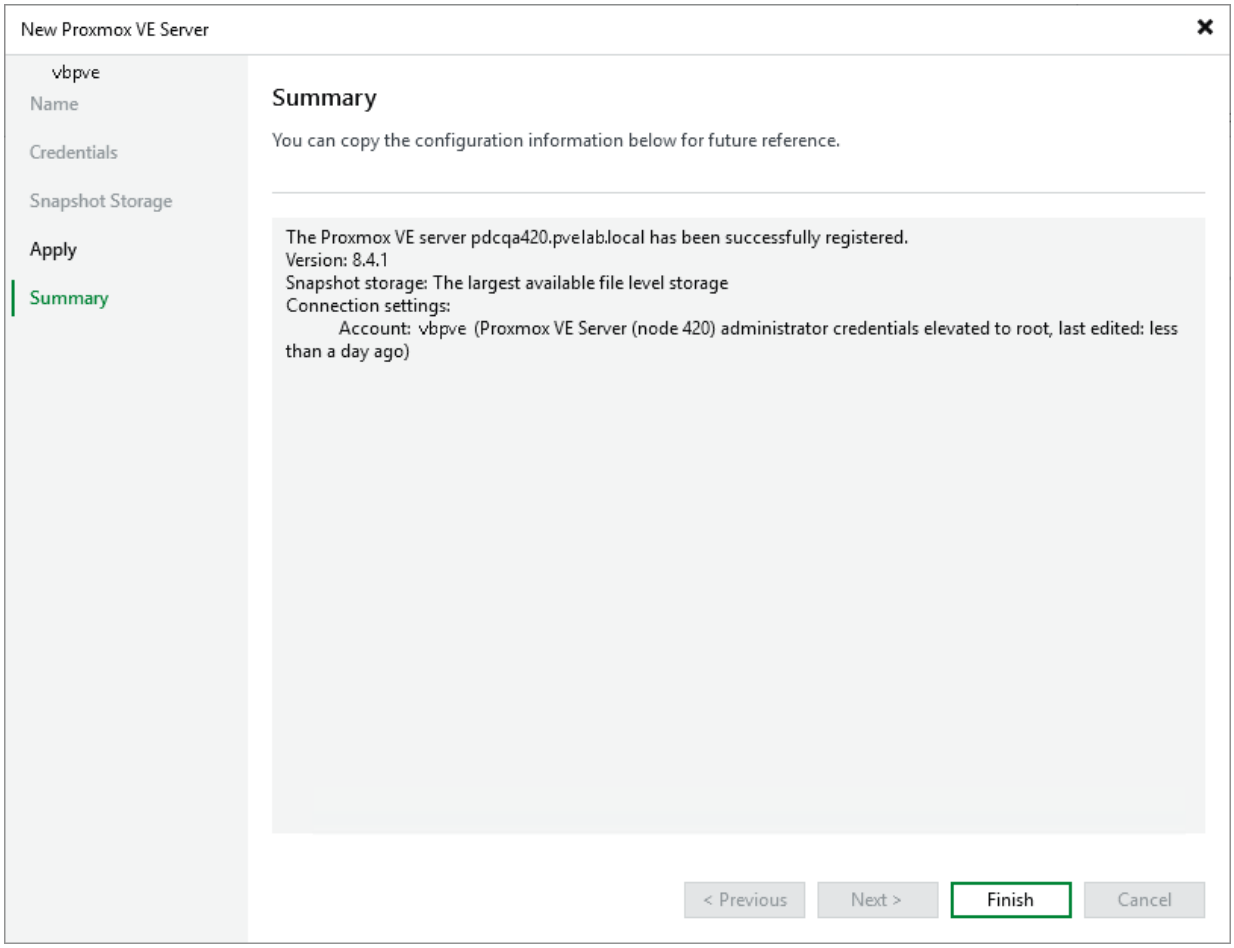

Editing Proxmox VE Server Properties

To edit properties of the Proxmox VE server added to the backup infrastructure, do the following:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select **Managed Servers > Proxmox VE > Clusters > pve-cluster**.
3. In the working area, select the Proxmox VE server and click **Edit Server** on the ribbon, or right-click the Proxmox VE server and select **Properties**.
4. Complete the **Edit Proxmox VE Server** wizard as described in section Adding Proxmox VE server to Backup Infrastructure.

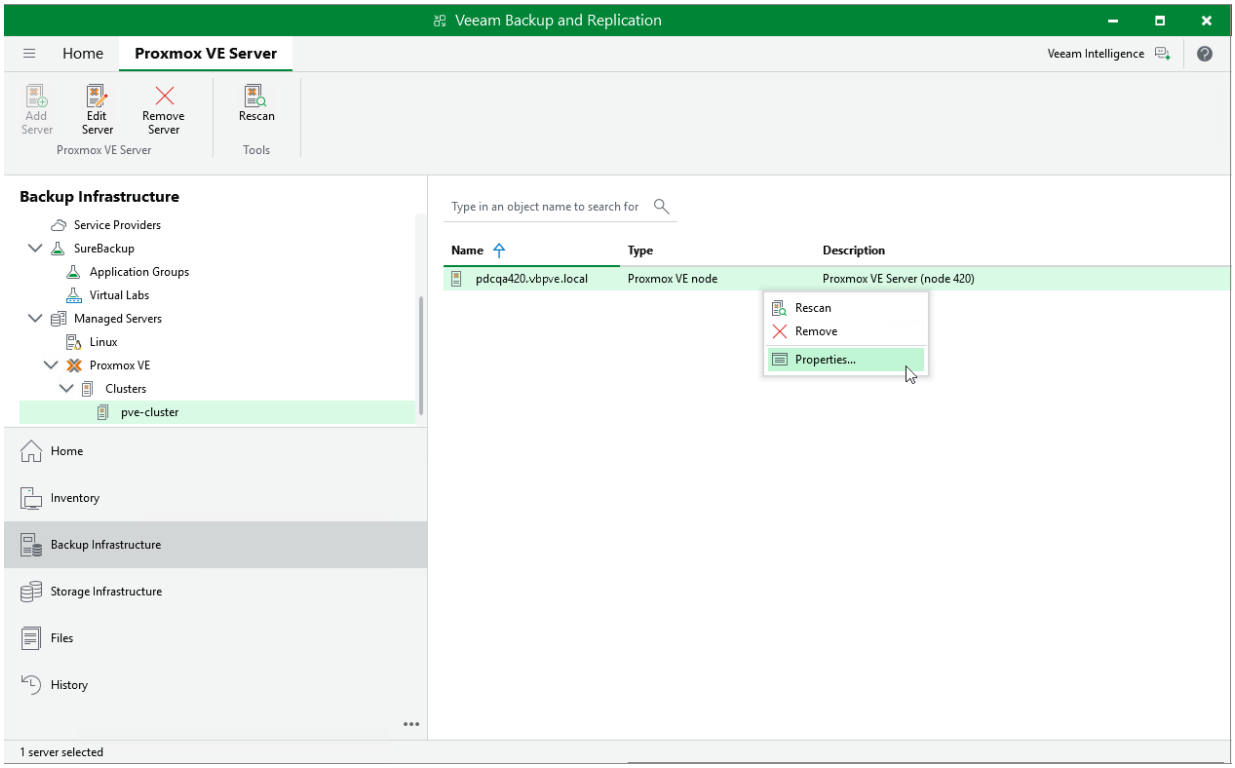

Rescanning Proxmox VE Server

Veeam Backup & Replication retrieves information about the Proxmox VE environment from the Proxmox VE server. However, the data synchronization process may take some time to complete. If you make any changes to the Proxmox VE environment and want the Veeam Backup & Replication console to display the changes immediately, you can rescan the Proxmox VE server manually.

To rescan the Proxmox VE server, do the following:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select **Managed Servers > Proxmox VE > Clusters > pve-cluster**.
3. In the working area, select the Proxmox VE server and click **Rescan** on the ribbon, or right-click the Proxmox VE server and select **Rescan**.

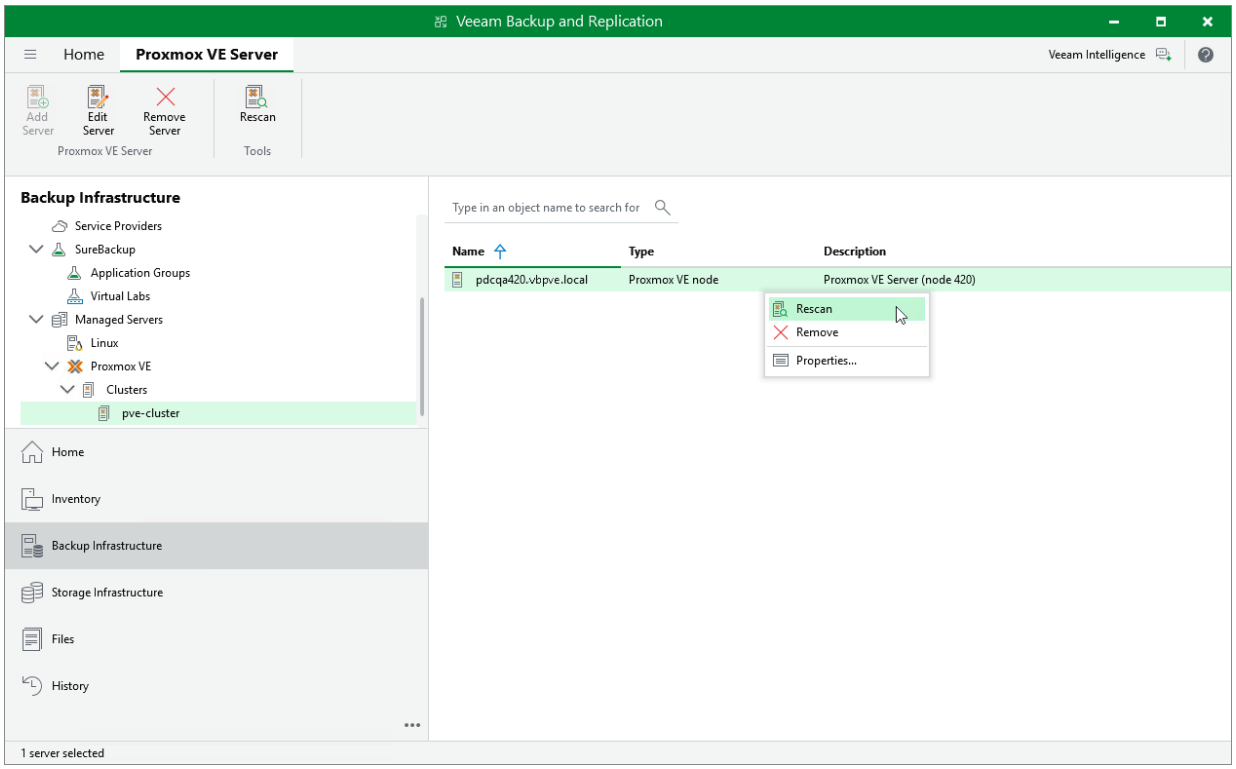

Removing Proxmox VE Server

If you do not want to protect resources managed by the connected Proxmox VE server anymore, you can remove it from the backup infrastructure.

To remove the Proxmox VE server from the backup infrastructure:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select **Managed Servers > Proxmox VE > Clusters > pve-cluster**.
3. In the working area, select the Proxmox VE server and click **Remove Server** on the ribbon, or right-click the Proxmox VE server and select **Remove**.

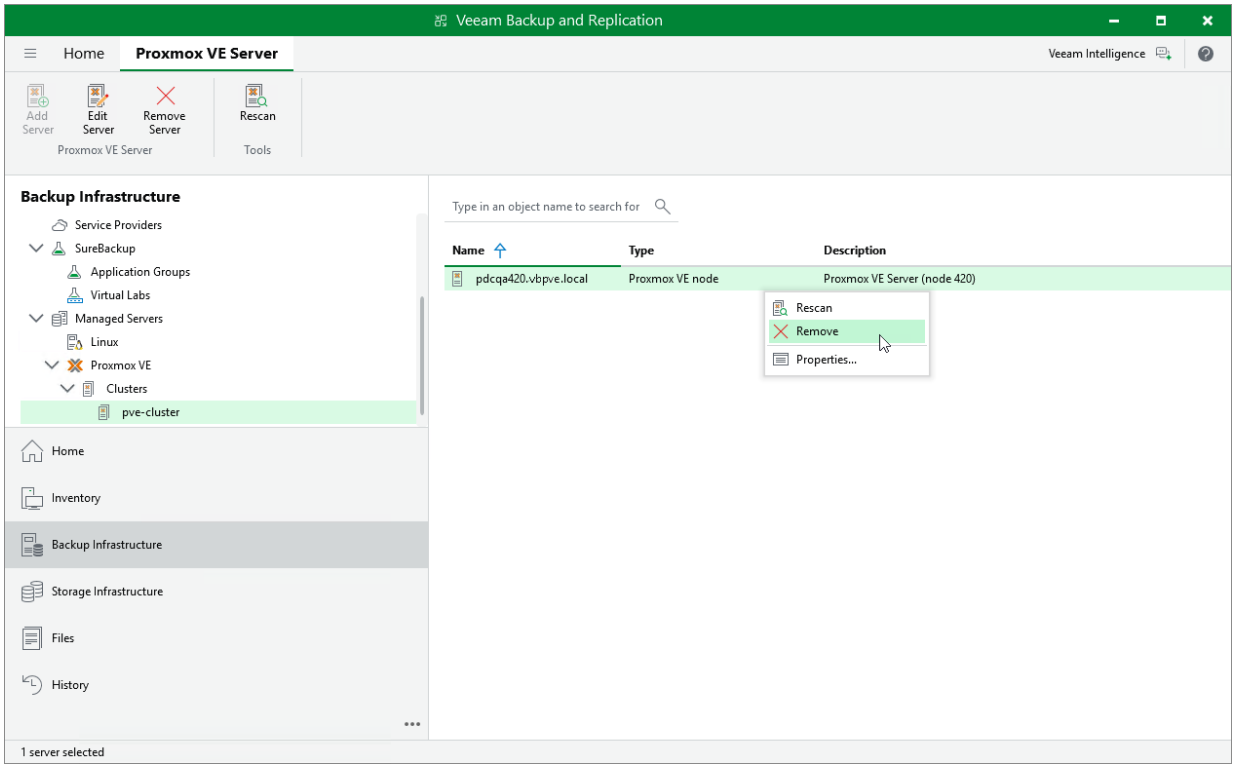

Managing Workers

To perform most data protection and disaster recovery operations, Veeam Backup & Replication uses workers. Workers are Linux-based VMs that process backup workload and distribute backup traffic when transferring data to backup repositories. Each worker is launched on a specific host for the duration of a backup or restore operation. As soon as a backup or restore session starts, Veeam Backup & Replication launches a worker, tests its configuration and installs system updates (if available). When the backup or restore session completes, Veeam Backup & Replication shuts down the worker VM so that it can be used for other sessions later.

> **IMPORTANT**
>
> To modify the worker settings, use the Veeam Backup & Replication console as described in section Editing Workers. Making any configuration changes to VMs running as workers manually in the Proxmox VE administration portal may cause technical issues.

Worker Lifecycle

When you add a worker to the backup infrastructure, its configuration is saved to the Veeam Backup & Replication configuration database, but no VM is actually deployed on the host unless you choose to test the configuration. In the latter case, a VM (worker VM) is deployed and shut down after the test operation completes.

As soon as a backup or restore session starts, Veeam Backup & Replication tries to launch the worker and test its configuration. If no worker VM has been previously deployed, Veeam Backup & Replication deploys the VM using the worker configuration saved to the configuration database. Then, Veeam Plug-in for Proxmox VE powers on the worker VM and installs system updates (if available). When the backup or restore session completes, Veeam Backup & Replication shuts down the worker VM so that it can be used for other sessions later.

During the lifecycle, a worker can obtain one of the following statuses:

- **Configured** – the worker configuration is added to the Veeam Backup & Replication configuration database.
- **Testing** – the worker configuration is being updated and tested.
- **Working** – the worker is processing a backup or restore operation.
- **Shut Down** – the worker is powered off.

Adding Workers

To deploy a worker and add it to the backup infrastructure, do the following:

1. Check prerequisites and limitations.
2. Launch the New Proxmox VE Worker wizard.
3. Specify worker VM configuration.
4. Specify worker network settings.
5. Finish working with wizard.

Before You Begin

Before you add a worker to the backup infrastructure, consider the following:

- It is recommended that workers are deployed on each node registered with a Proxmox VE cluster. If no worker is deployed on the node, performance of backup and restore operations will be affected as Veeam Backup & Replication will use a worker deployed on another node.
- Each worker must be provided with sufficient compute resources to handle backup and restore tasks in parallel. The maximum number of concurrent tasks is configured in worker settings — if this number is exceeded, the worker will not start a new task until one of the current tasks finishes.
- You can change the maximum number of concurrent tasks (the best practice is to allocate 1 vCPU and 1 GB RAM for each additional task) while deploying a new worker or editing settings of an existing one.

Step 1. Launch New Proxmox Worker

To launch the **New Proxmox VE Worker** wizard, do the following:

1. In the Veeam Backup & Replication console, open the **Backup Infrastructure** view.
2. In the inventory pane, select **Backup Proxies**.
3. On the ribbon, select **Add Proxy**.
4. Click **Proxmox VE worker**.

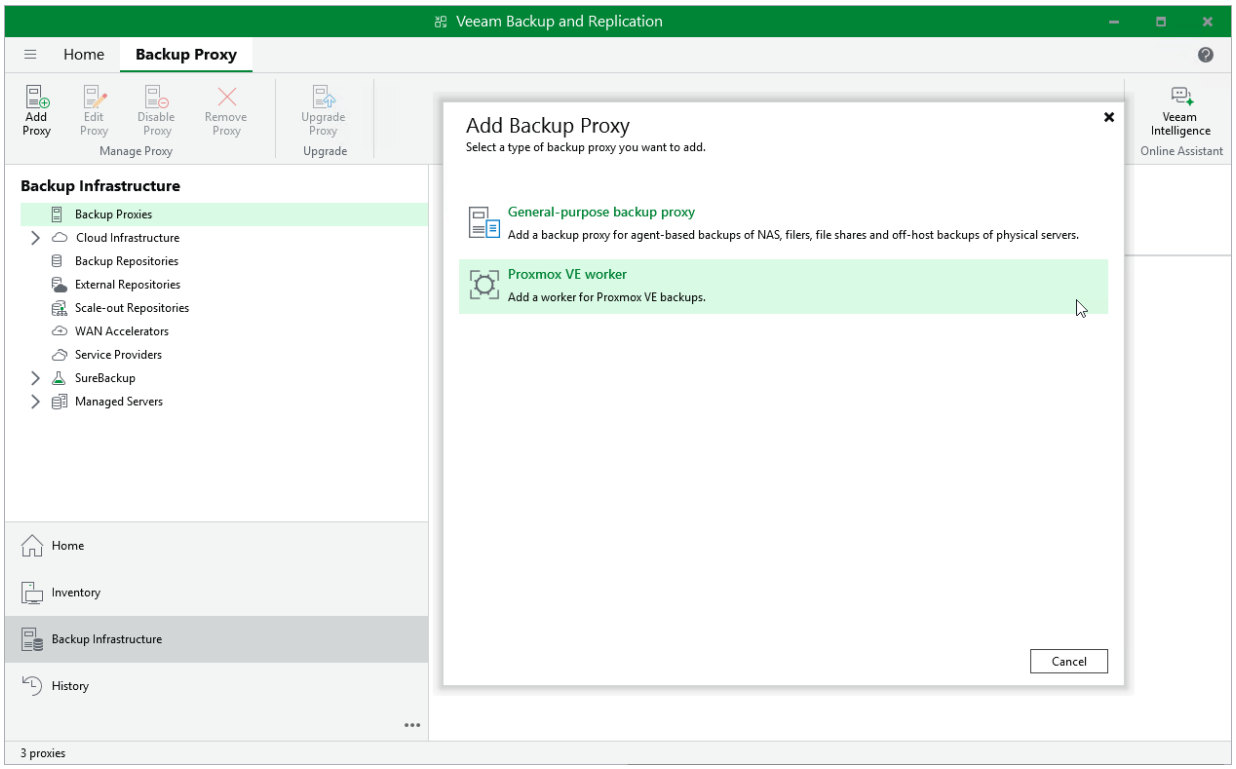

Step 2. Specify Worker VM Settings

At the **Virtual Machine** step of the wizard, do the following:

1. Click **Choose** next to the **Host** field to specify a host where the worker will be launched.

   Make sure that the default *local* storage is enabled on the selected host. If you cannot use the default storage in your environment, contact Veeam Customer Support.

2. In the **Name** field, specify a name for the worker. The maximum length of the name is 40 characters; the following characters are only supported: a-z, A-Z, 0-9, -.

3. Click **Choose** next to the **Storage** field to select storage where system files of the worker will be stored. For storage to be displayed in the list of available storage, it must be configured in the virtual environment as described in Proxmox VE documentation.

   Make sure that the selected storage supports snapshots.

4. In the **Worker description** field, provide a description for future reference. The maximum length of the description is 1024 characters.

5. In the **Max concurrent tasks** field, specify the number of tasks that the worker will be able to handle in parallel. If this value is exceeded, the worker will not start processing a new task until one of the currently running tasks finishes.

   The default number of concurrent tasks is set to *4*. When you change this value, the wizard automatically adjusts the amount of resources that will be allocated to the worker. If you want to specify the amount of resources manually, click **Advanced proxy settings**.

NOTE

- When performing data protection and disaster recovery operations, Veeam Backup & Replication initiates a new task for each VM that is being processed.
- When processing VMs, Veeam Backup & Replication adjusts the number of concurrent tasks taking into account a specific limit of backup operations (set to 4 by default) that applies to storage in order to avoid excessive load on the production environment. For example, you configure a worker to process maximum 10 VMs simultaneously, while 5 of these VMs store their files on one storage and the other 5 VMs — on another storage; in this case, the worker will process 8 VMs simultaneously — 4 VMs for each storage.

  The default backup operation limit cannot be changed using the Veeam Backup & Replication console. To change the limit, contact Veeam Customer Support.

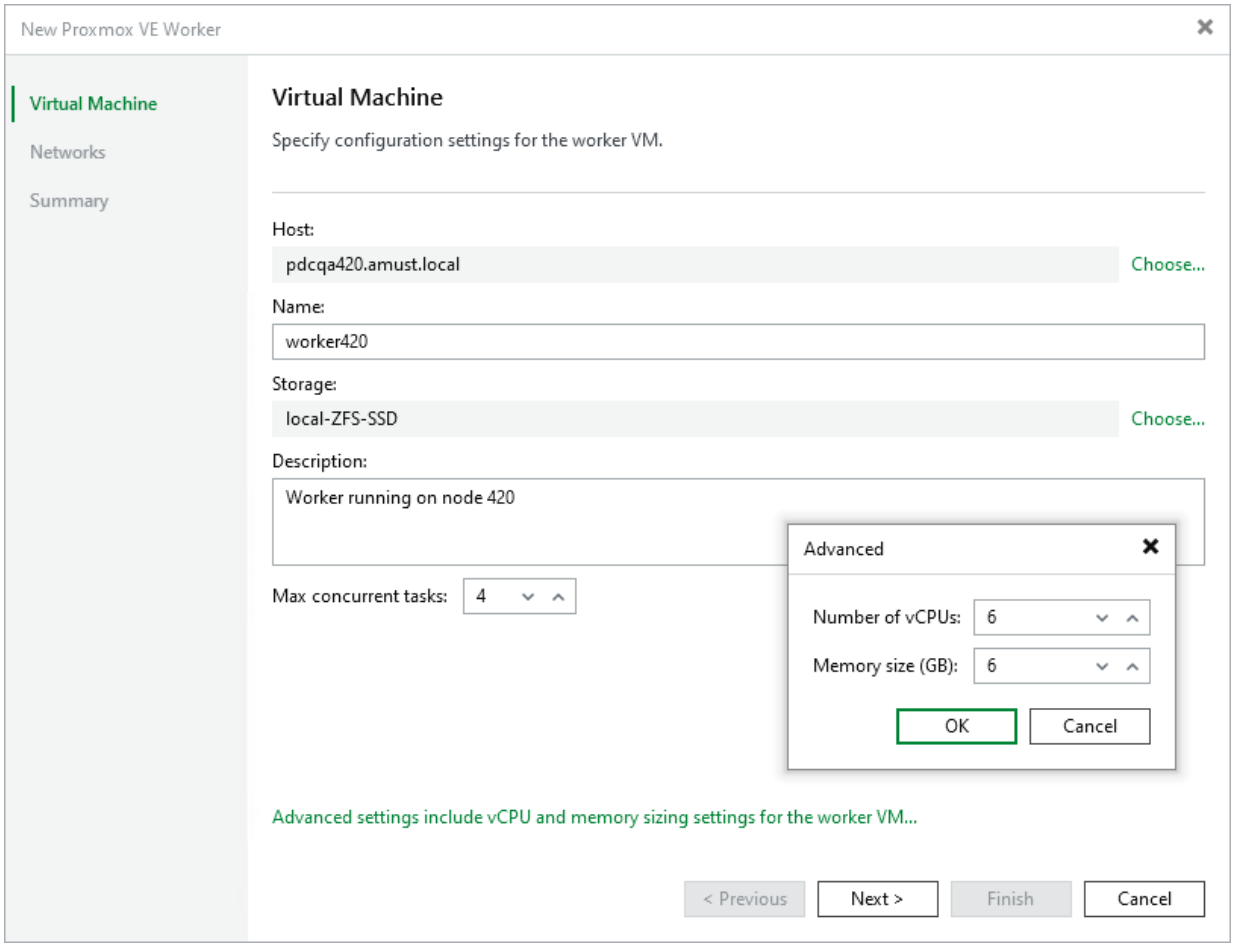
New Proxmox VE Worker
Virtual Machine
Networks
Summary
Virtual Machine
Specify configuration settings for the worker VM.
Host:
pdcqa420.amust.local
Choose...
Name:
worker420
Storage:
local-ZFS-SSD
Choose...
Description:
Worker running on node 420
Max concurrent tasks:
4
Advanced
Number of vCPUs:
6
Memory size (GB):
6
OK
Cancel
Advanced settings include vCPU and memory sizing settings for the worker VM...
< Previous
Next >
Finish
Cancel

Step 3. Configure Network Settings

At the **Networks** step of the wizard, do the following:

1. Click **Add** to configure worker network interfaces:

   a. From the **Network** drop-down list, select a network to which the worker network interface will be connected.

      For a network to be displayed in the list of available networks, it must be configured in the virtual environment as described in Proxmox VE documentation.

   b. In the **Description** field, provide a network interface description for future reference.

   c. If DHCP is enabled in the selected network, the IP address of the worker can be obtained automatically.

      If DHCP is disabled in the selected network, or you want to specify an IP address, select the **Use the following IP address** option and enter the worker IP address, subnet mask and default gateway.

   To add more network interfaces, repeat the step and specify the network order using the **Up** and **Down** buttons. For more information on multi-network configuration, see section Appendix. Configuring Multiple Networks.

2. If DHCP is enabled in any network to which the worker will be connected, DNS settings of the worker can be obtained automatically. To configure DNS settings manually, click **Obtain automatically** and do the following in the **DNS Server Settings** window:

   a. Select the **Use the following DNS server address** option.

   b. Enter the IP addresses of the preferred and alternate DNS servers.

   c. Click **OK**.

> **NOTE**
>
> Since workers are Linux-based VMs, they have the same limitations that apply to machines running the Rocky Linux operating system. Therefore, DNS settings cannot be configured separately for each network added to the worker.

3. To check for available package updates for the worker, Veeam Backup & Replication automatically connects to Veeam repositories over the internet. If the worker is not connected to the internet, you can instruct Veeam Backup & Replication to use an HTTP proxy that will provide access to the necessary repositories. To specify HTTP proxy settings, click **Advanced** and do the following in the **Advanced Settings** window:

   a. Select the **Check for updates online** check box.

   b. Select the **Use the following internet proxy settings** check box.

   c. In the **Host** field, enter the IP address or FQDN of the web proxy.

   d. In the **Port** field, enter the port used on the web proxy for HTTP or HTTPS connections.

e. [Applies only if the HTTP proxy requires authentication] Select the **Use authentication** check box and select credentials of the account configured on the HTTP proxy to access the internet.

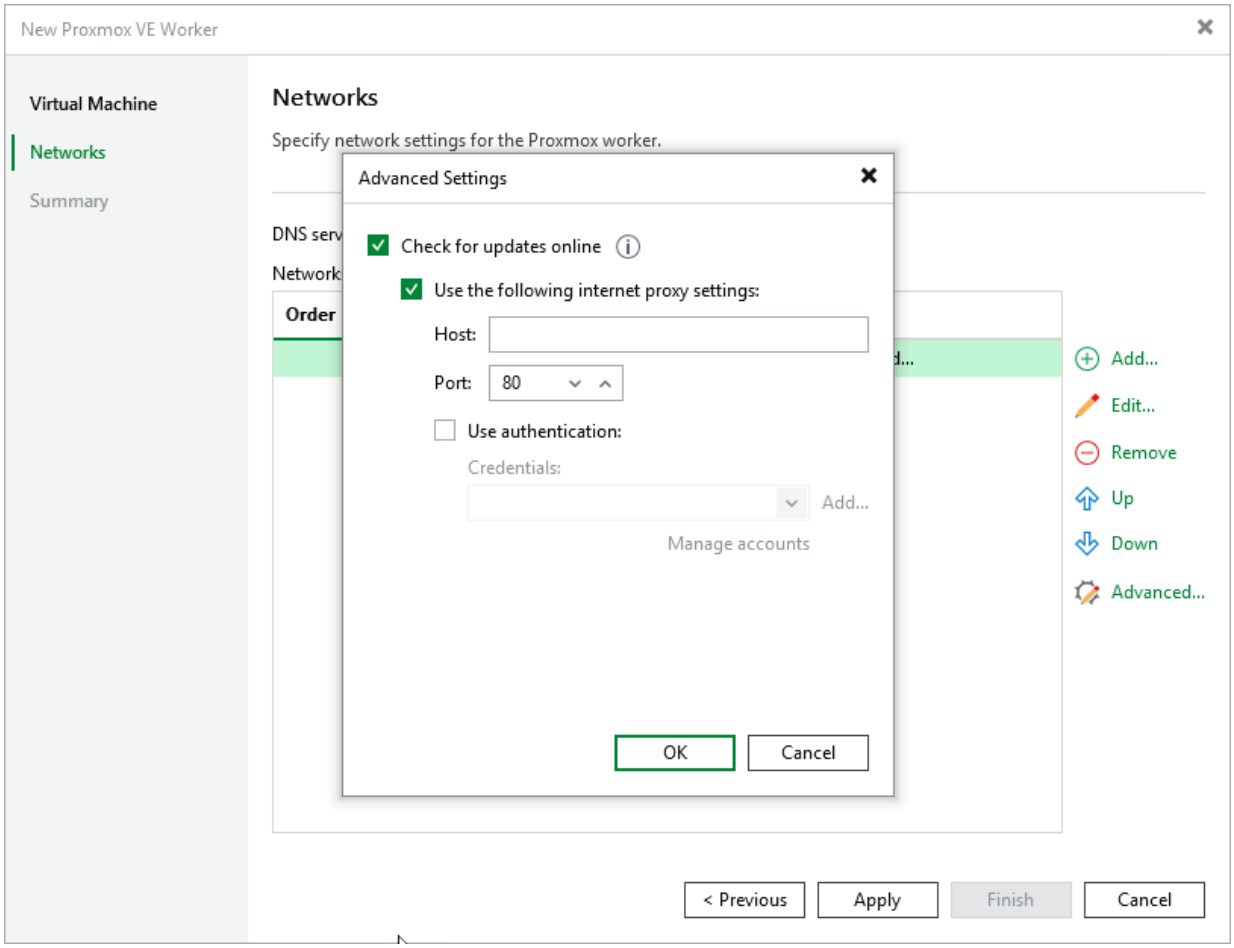

Step 4. Finish Working with Wizard

At the **Summary** step of the wizard, review summary information and click **Finish**.

> **TIP**
>
> To test the worker, select the **Test the worker configuration when I click Finish** check box and then click **Finish**.

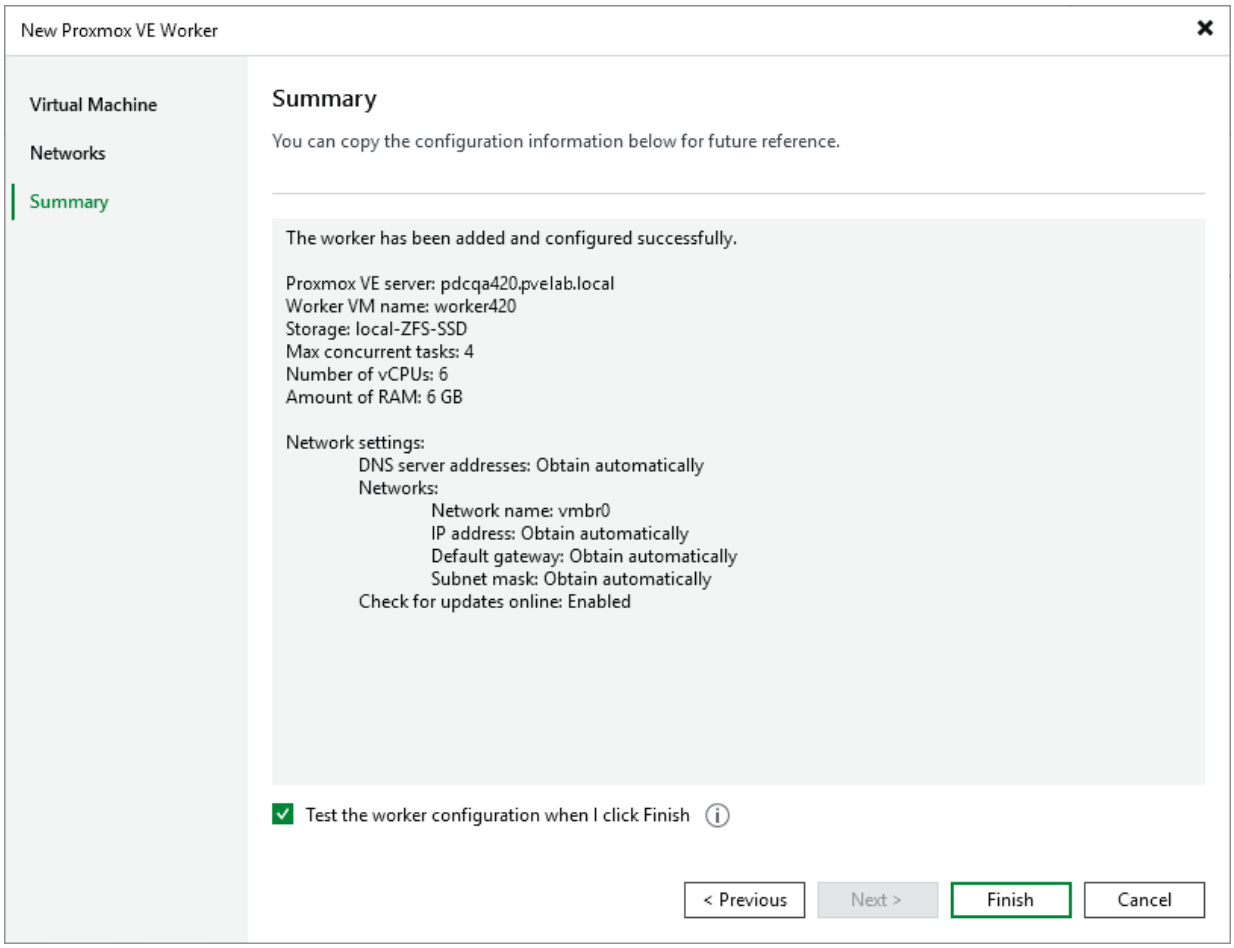

Testing Workers

Before using a dedicated worker for a backup or restore operation, Veeam Backup & Replication automatically tests its configuration — verifies that the worker service can start successfully, checks that the worker can connect to the backup server and to the host, and installs available updates.

If you want to ensure that the worker configuration is correct before it is used for a backup or restore operation, you can start a worker configuration test manually:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select **Backup Proxies**.
3. In the working area, select the necessary worker and click **Test Worker** on the ribbon.

   Alternatively, right-click the worker and select **Test**.

> **TIP**
>
> As soon as Veeam Backup & Replication finishes the worker configuration test, the worker will be powered off. You can review details of the test session in system logs as described in the Veeam Backup & Replication User Guide, section Viewing History Statistics.

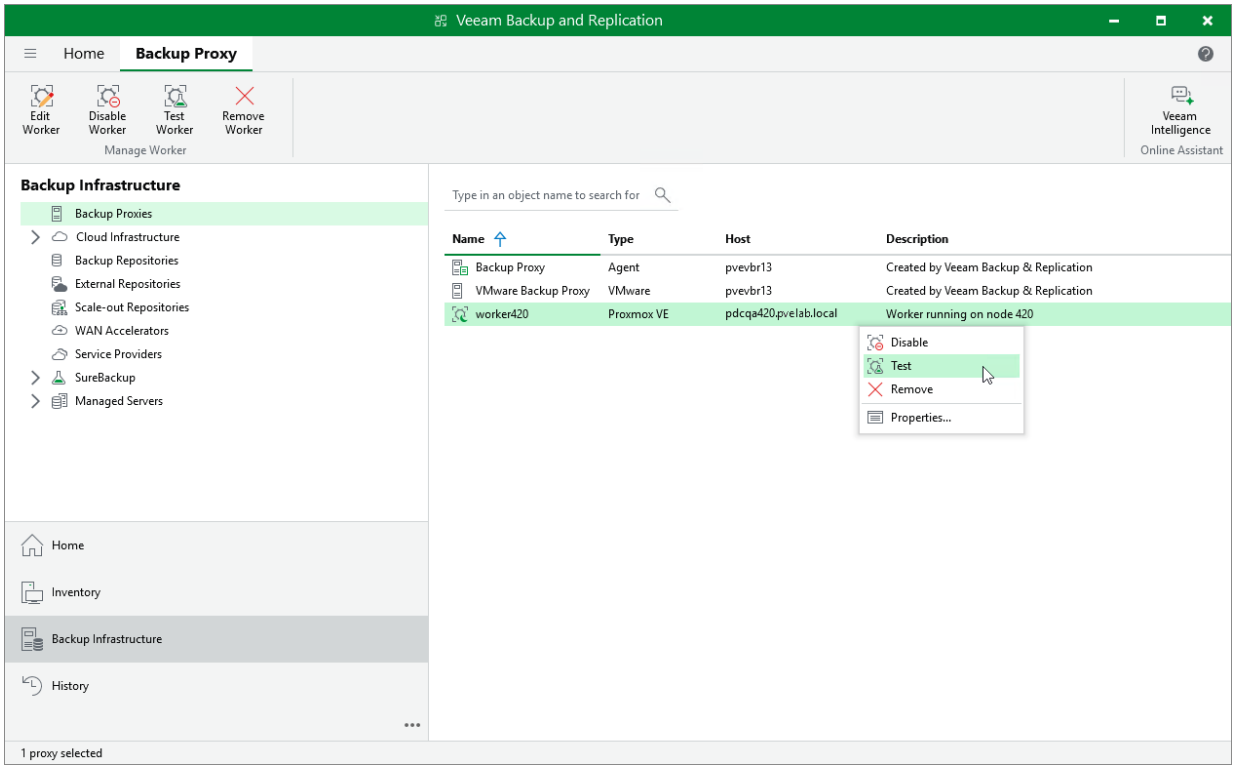

Enabling and Disabling Workers

By default, workers are launched when jobs or restore sessions start. However, you can temporarily disable a worker — this may be helpful when you reconfigure a worker and you do not want it to be used for a backup or restore operation. You will still be able to enable the disabled worker at any time you need.

To enable or disable a worker, do the following:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select **Backup Proxies**.
3. In the working area, select the necessary worker and click **Disable Worker** or **Enable Worker** on the ribbon.

   Alternatively, right-click the worker and select **Disable** or **Enable**.

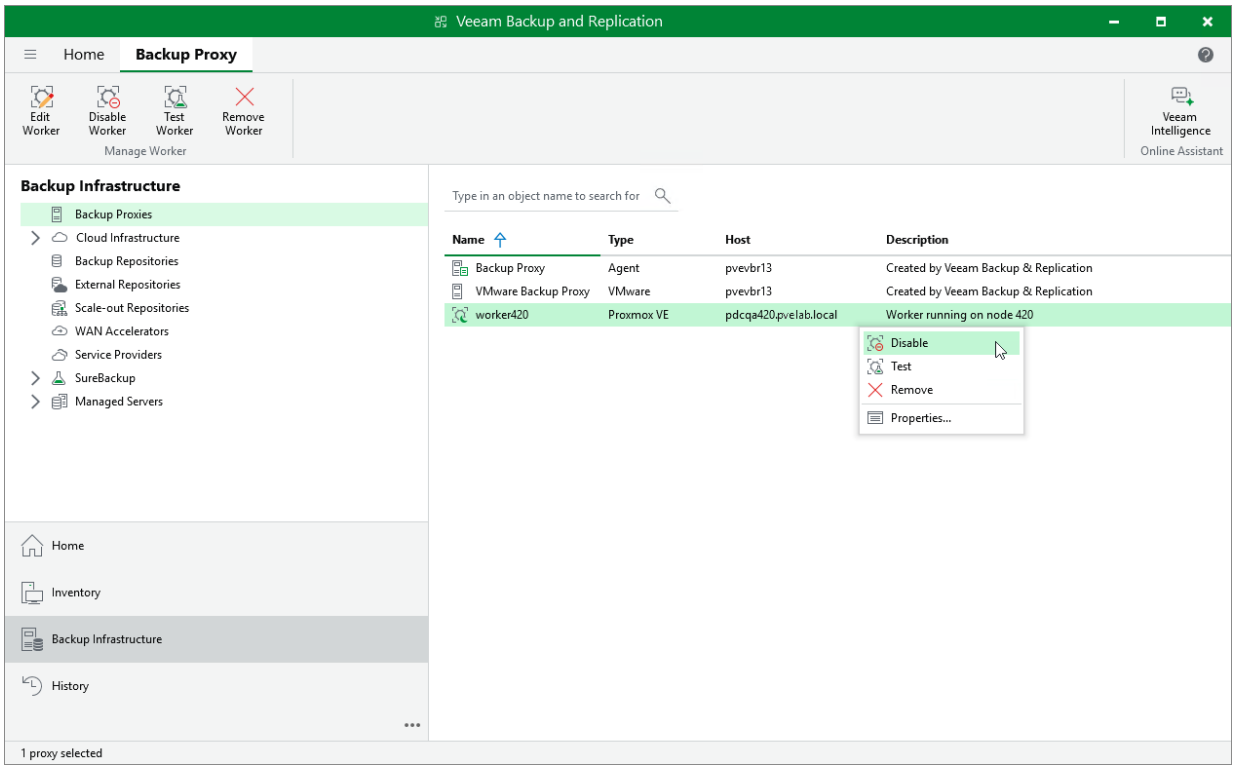

Editing Workers

For each worker, you can modify settings specified while adding the worker to the backup infrastructure:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select **Backup Proxies**.
3. In the working area, select the necessary worker and click **Edit Worker** on the ribbon.

   Alternatively, right-click the worker and select **Properties**.
4. Complete the **Edit Proxmox VE Worker** wizard:

   a. To provide a new name and description for the worker or to modify the number of tasks that the worker is able to handle in parallel, follow the instructions provided in section Adding Workers (step 2).

   b. To change the network to which the worker is connected or to specify a new IP address for the worker, follow the instructions provided in section Adding Workers (step 3).

   c. To save changes made to the worker settings, click **Finish**.

> **IMPORTANT**
>
> It is not recommended that you decrease the amount of allocated resources or modify the network settings while the worker is currently transferring data. In this case, Veeam Backup & Replication will terminate the related operations, power off the worker and update the settings immediately.

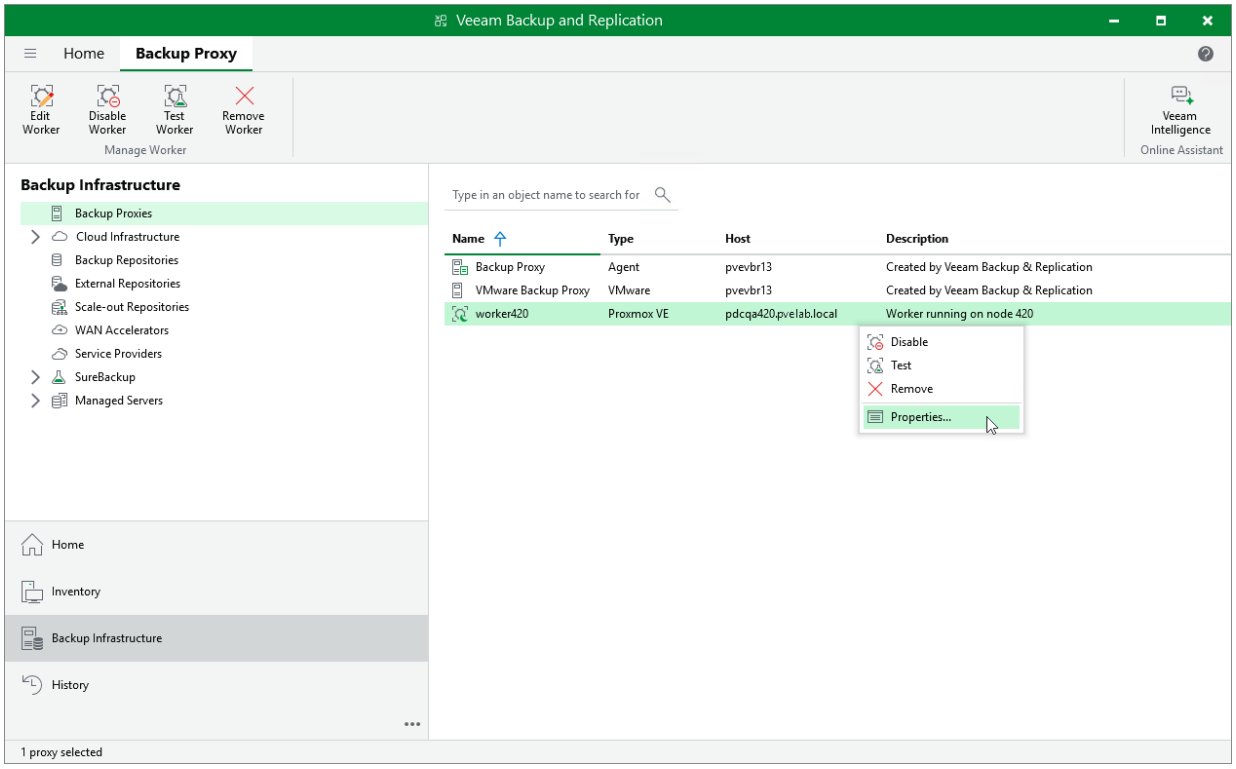

Disabling Automatic Worker Updates

When launching a worker for a backup or restore operation, Veeam Backup & Replication automatically downloads updates from Veeam repositories and installs them on the worker. If the worker is not connected to the internet, you can instruct Veeam Backup & Replication to use an HTTP proxy that will provide access to the necessary repositories.

If a worker does not have access to the internet and no HTTP proxy is configured for the worker, you can disable automatic updates to avoid connection failures and eliminate session warnings:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select **Backup Proxies**.
3. In the working area, select the necessary worker and click **Edit Worker** on the ribbon.

   Alternatively, right-click the worker and select **Properties**.
4. At the **Networks** step of the **Edit Proxmox VE Worker** wizard, click **Advanced** and clear the **Check for updates online** check box. Then, click **Finish** to save changes made to the worker settings.

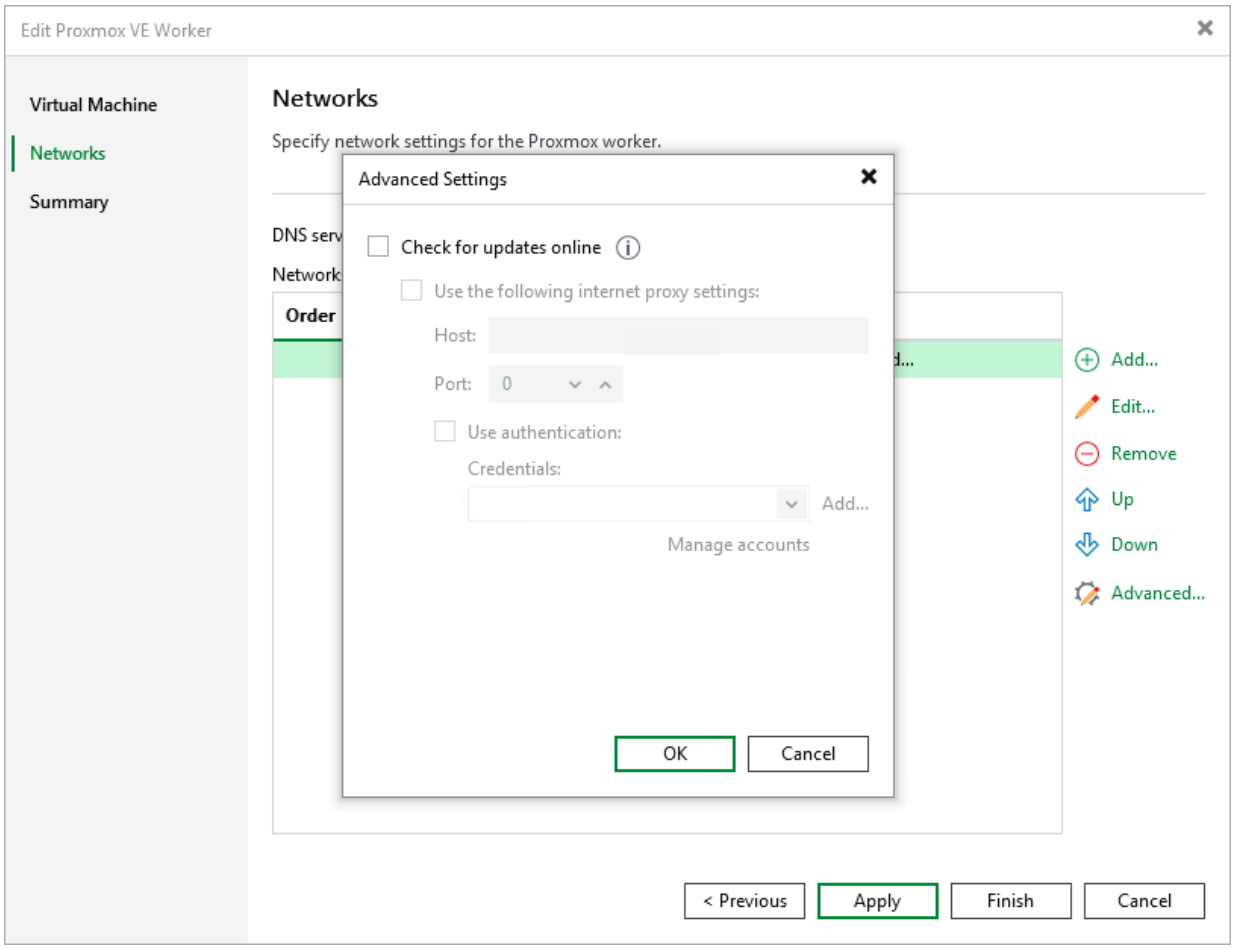

Removing Workers

Veeam Backup & Replication allows you to permanently remove workers if you no longer need them. Note that you can remove a worker only when it is not processing a backup or restore operation.

To remove a worker, do the following:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select **Backup Proxies**.
3. In the working area, select the necessary worker and click **Remove Worker** on the ribbon.

   Alternatively, right-click the worker and select **Remove**.
4. In the **Veeam Backup & Replication** window, confirm that you want to permanently delete the worker.

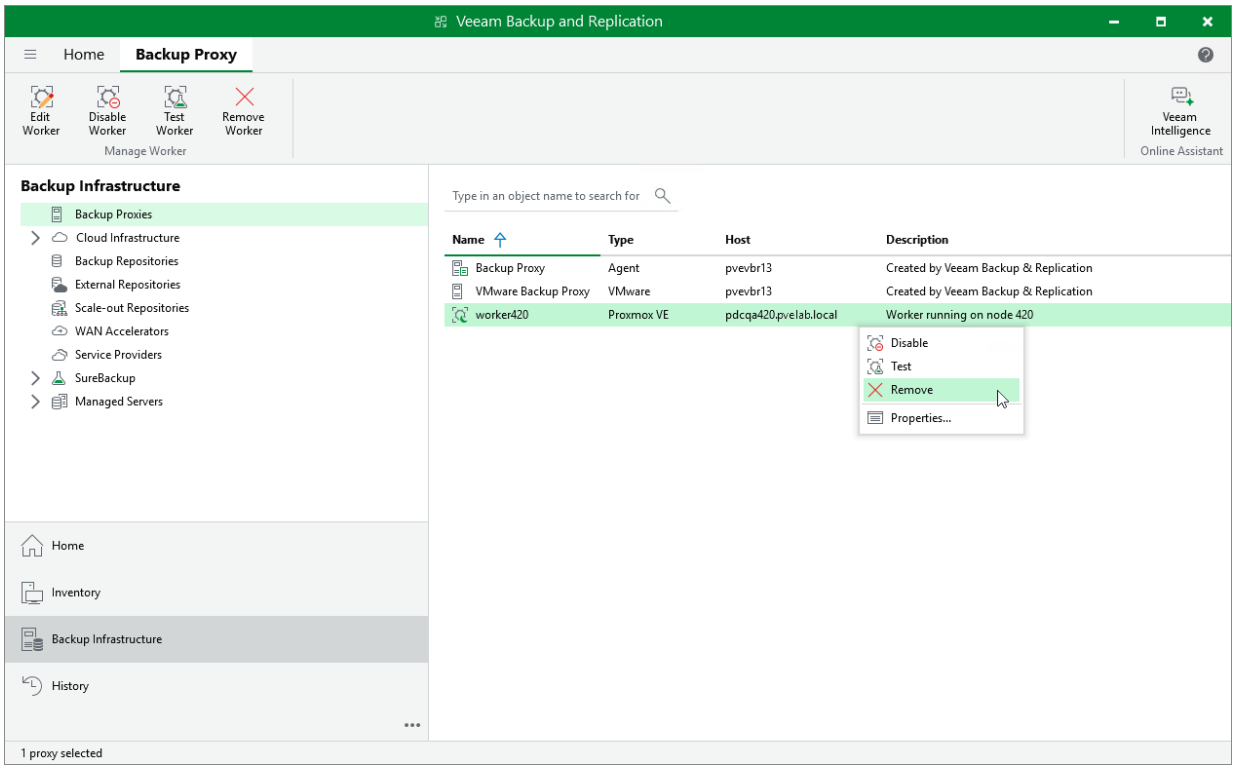

Configuring General Settings

Veeam Backup & Replication allows you to configure general settings that are applied to all performed operations and deployed architecture components:

- Configure email settings for automated delivery of reports.
- Configure notification settings.

Configuring Email Settings

You can specify email notification settings for automated delivery of job results. To connect an SMTP server that will be used for sending email notifications:

1. From the main menu of the Veeam Backup & Replication console, select **Options**.
2. Switch to the **E-mail Settings** tab.
3. Select the **Enable e-mail notifications** check box.
4. Configure mail server settings.
6. In the **From** field, enter an email address of the notification sender. This email address will be displayed in the **From** field of notifications.
7. In the **To** field, enter an email address of a recipient. Use a semicolon to separate multiple recipient addresses.
8. In the **Subject** field, specify a subject for notifications. You can use the following runtime variables:
   - *%JobName%* — a job name.
   - *%JobResult%* — a job result.
   - *%ObjectCount%* — the number of VMs in a job.
9. Choose whether you want to receive email notifications in case jobs complete successfully, complete with warnings or complete with errors.
10. Select the **Suppress notifications until the last job retry** check box to receive a notification about the final job status. If you do not select this check box, the Veeam Backup & Replication will send one notification for every job retry.
11. Click **Apply**.

> **TIP**
>
> Veeam Backup & Replication allows you to send a test message to check whether you have configured the settings correctly. To do that, click **Test Message**. A test message will be sent to the specified email address.

Configuring Mail Server Settings

To configure mail server settings, choose whether you want to employ SMTP server, Google Gmail or Microsoft 365 authentication for your mail server.

Using SMTP Server Basic Authentication

To employ the SMTP server basic authentication to connect to your mail server, do the following in the **Options** window:

1. From the **Mail server** drop-down list, select **SMTP server (basic authentication)**.
2. In the **SMTP server** field, enter a DNS name or an IP address of the SMTP server. All email notifications (including test messages) will be sent by this SMTP server.

3. Click **Advanced** next to the **Mail server** field and configure SMTP server settings:
   a. In the **Port** field, specify a communication port for SMTP traffic. The default SMTP port is 587.
   b. In the **Timeout** field, specify a connection timeout for responses from the SMTP server.
   c. For an SMTP server with SSL/TLS support, select the **Connect using SSL** check box to enable SSL data encryption.
   d. If your SMTP server requires authentication, select the **This SMTP server requires authentication** check box and specify credentials that will be used to connect to the SMTP server.

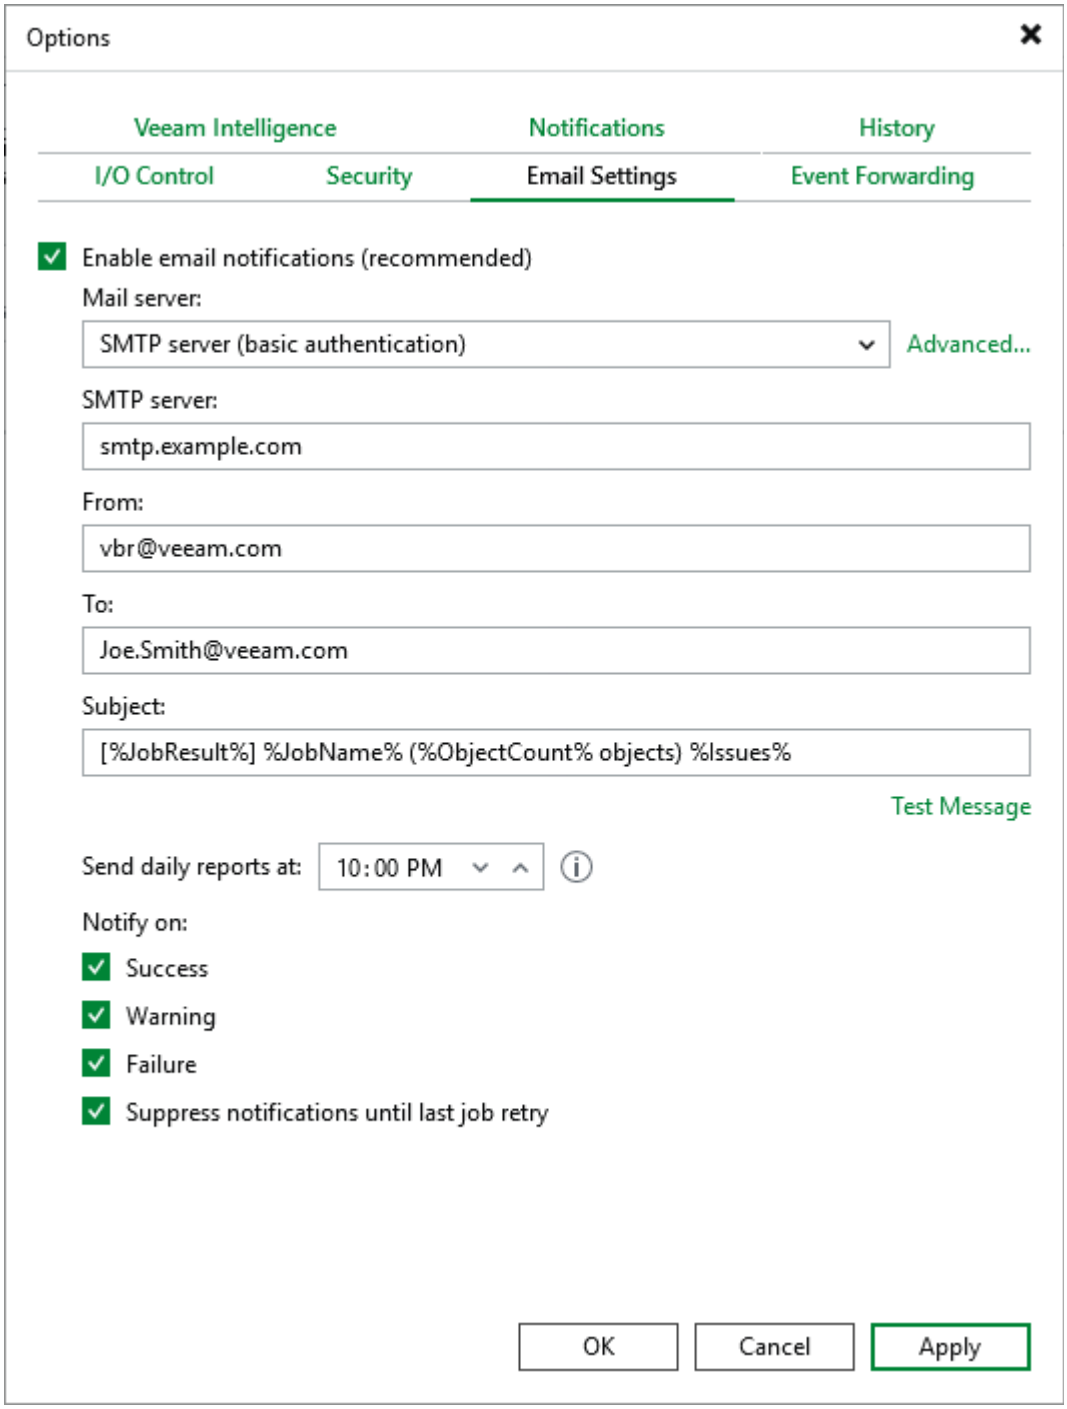

Using Google Gmail Modern Authentication

To employ the Google Gmail modern authentication to connect to your mail server, do the following in the **Options** window:

1. From the **Mail server** drop-down list, select *Google Gmail (modern authentication)*.
2. Click **Sign in with Google**. You will be redirected to the authorization page.
3. On the authorization page, specify a Google account to connect to the Veeam Backup & Replication application. Note that you must also select the **Send email on your behalf** check box.

> **TIP**
>
> If you want to use your own web application for email notifications, do the following:
>
> 1. Register a new client application in the Google Cloud console for Veeam Backup & Replication to be able to use OAuth 2.0 to access Google Cloud APIs. When registering the application, it is recommended to use a dedicated service account with granular *SendMail* permissions.
> 2. In the **Options** window, click **Advanced**.
> 3. In the **Advanced** window, select the **Use custom registration settings** check box, and provide the application client ID and client secret created for the application as described in Google Cloud documentation.
> 4. Click **Sign in with Google**. You will be redirected to the authorization page.
> 5. On the authorization page, specify a Google account to connect to the registered application. Note that you must also select the **Send email on your behalf** check box.

If the authentication process completes successful, Veeam Backup & Replication will display a message notifying that the token is valid. If the token gets revoked or if the Google account password changes, click **Re-authorize** to update the configuration settings.

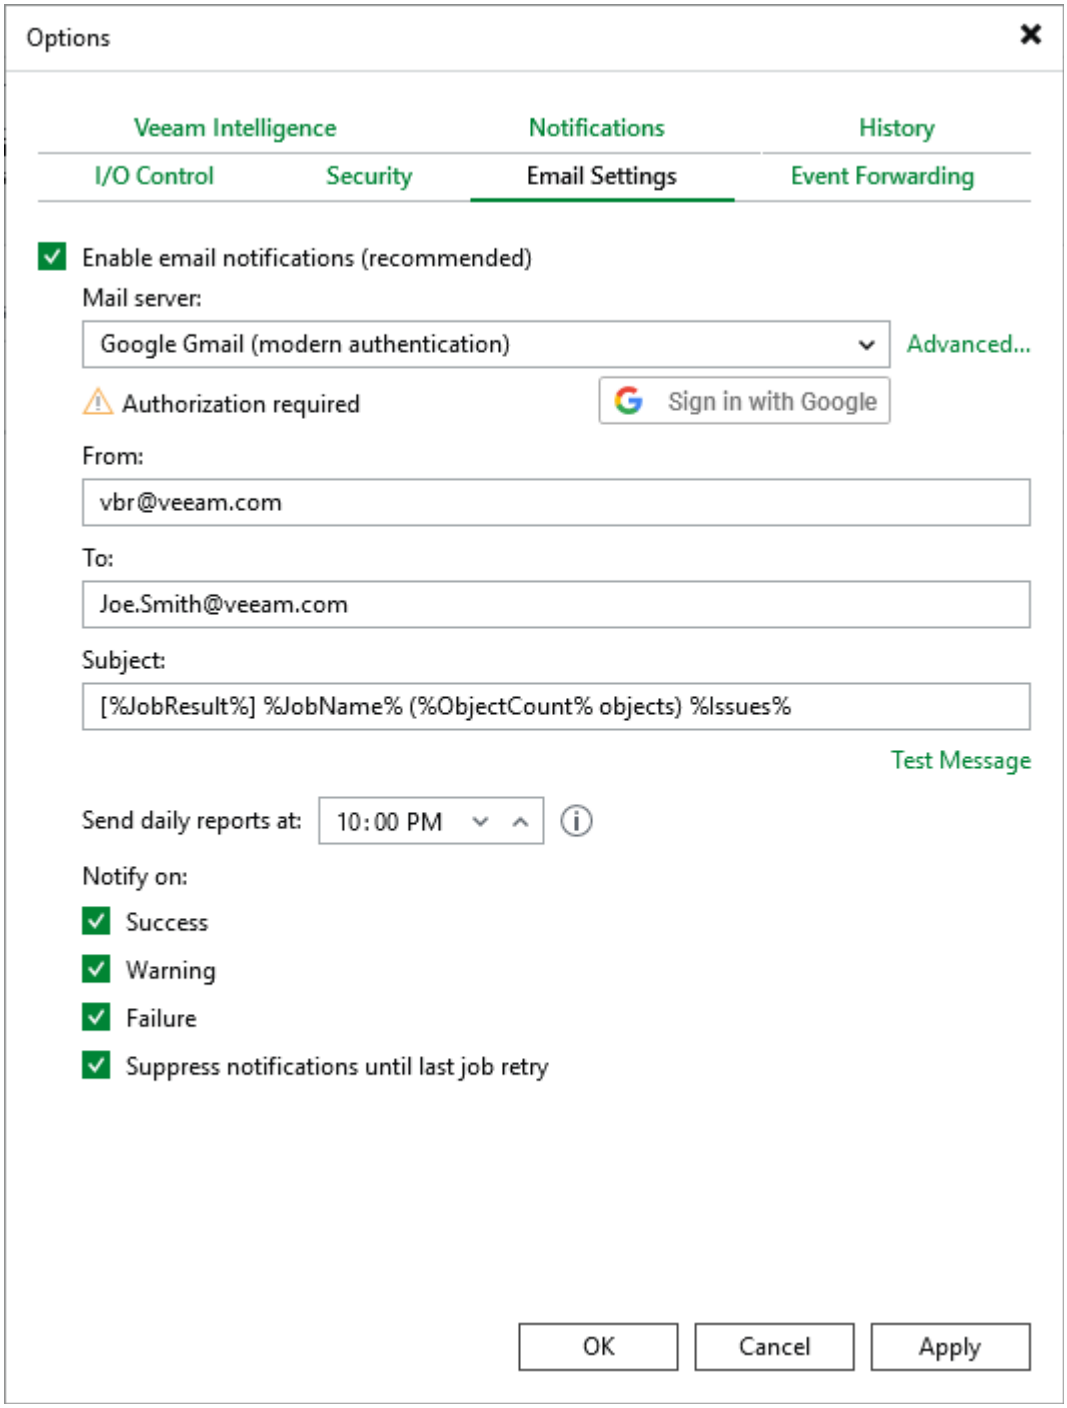

Using Microsoft 365 Modern Authentication

To employ the Microsoft 365 modern authentication to connect to your mail server, do the following in the **Options** window:

1. From the **Mail server** drop-down list, select *Microsoft 365 (modern authentication)*.
2. Click **Authorize now**. You will be redirected to the authorization page.

3. On the authorization page, specify an Exchange Online account to connect to the Veeam Backup & Replication application. Note that you must also select the **Consent on behalf of your organization** check box.

   To sign in with Exchange Online credentials, you may need to turn off the Internet Explorer Enhanced Security Configuration option in Server Manager as described in Microsoft documentation.

> **TIP**
>
> If you want to use your own web application for email notifications, do the following:
>
> 1. Register a new client application in the Microsoft Azure portal for Veeam Backup & Replication to be able to use OAuth 2.0 to access Microsoft Azure APIs. When registering the application, it is recommended to use a dedicated service account with granular *SendMail* permissions.
> 2. In the **Options** window, click **Advanced**.
> 3. In the **Advanced** window, select the **Use custom registration settings** check box, and provide the application client ID and tenant ID created for the application as described in Microsoft documentation.
> 4. Click **Authorize now**. You will be redirected to the authorization page.
> 5. On the authorization page, specify a Exchange Online account to connect to the registered application. Note that you must also select the **Send email on your behalf** check box.

If the authentication process completes successful, Veeam Backup & Replication will display a message notifying that the token is valid. If the token gets revoked or if the Microsoft account password changes, click **Re-authorize** to update the configuration settings.

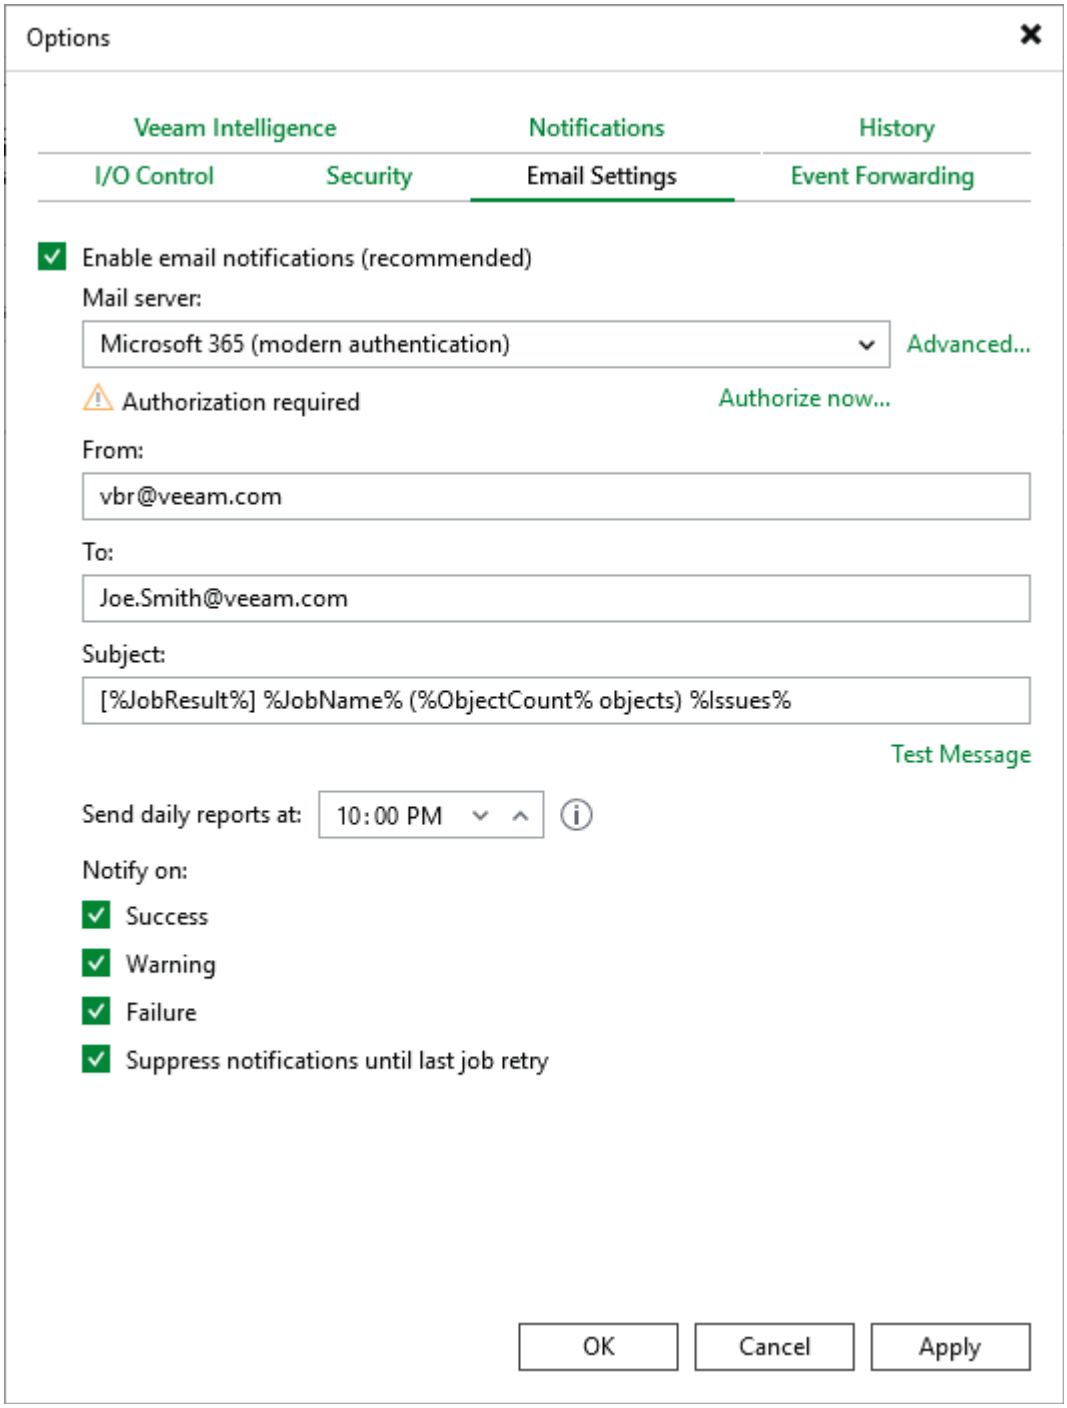

Configuring Notification Settings

You can enable notifications for Veeam Backup & Replication events that may require your actions:

1. From the main menu of the Veeam Backup & Replication console, select **Options**.
2. Switch to the **Notifications** tab.
3. In the **Backup storage** section, choose whether you want to receive notifications when backup repositories used as target locations for VM backups start running out of free space. While processing VMs included into backup jobs, Veeam Backup & Replication analyzes the amount of storage space left in target repositories and displays warnings in job session details if a specific threshold is breached.
4. In the **Production datastores** section, choose whether you want to receive notifications when Proxmox VE storage disks used as target locations for VM snapshots start running out of free space. While processing VMs included into backup jobs, Veeam Backup & Replication analyzes the amount of space left on target storage disks and displays warnings in job session details if a specific threshold is breached.

> **TIP**
>
> If Veeam Backup & Replication detects a target storage disk that is about to run out of free space while processing a VM, it will either skip the VM from processing or create a snapshot of the VM anyway, which may result in storage disruptions in the production environment. To avoid the latter, you can instruct Veeam Backup & Replication to skip VMs from processing if a specific threshold is breached.

5. In the **Support expiration** section, choose whether you want to receive notifications when the Production Support and Maintenance agreement included into your Subscription license is about to expire. When Veeam Backup & Replication detects that there are less than 14 days left before the support expiration date, it sends an email notification to the recipient specified in the general email settings.

   For more information on how to track the support expiration date, see the Veeam Backup & Replication User Guide, section Viewing License Information.

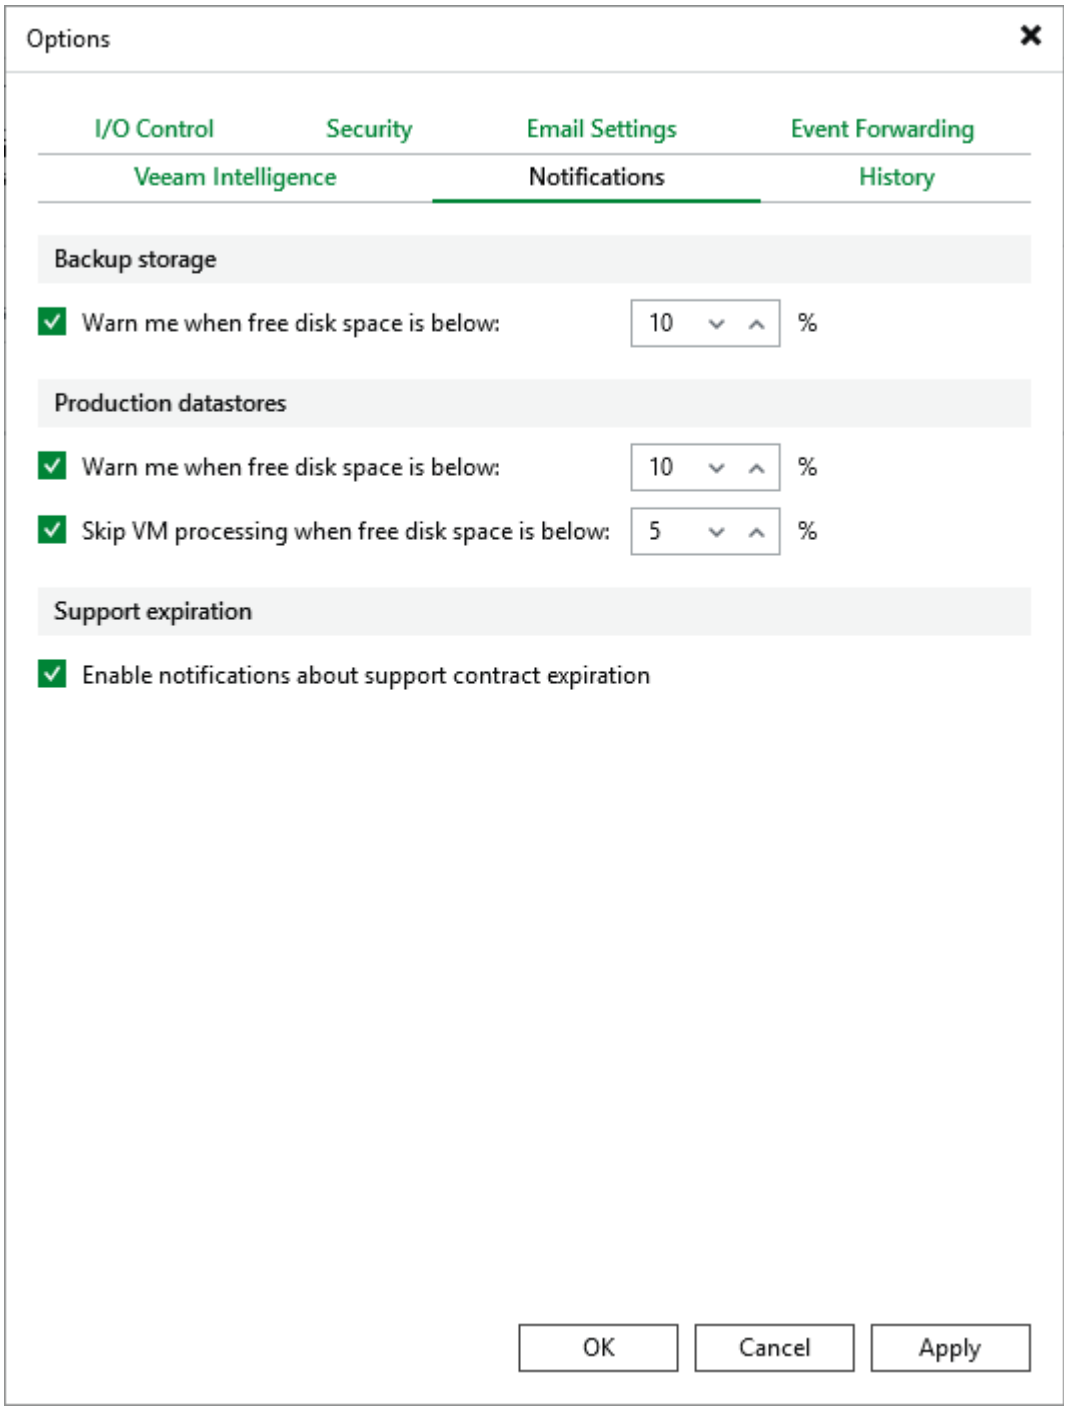

Performing Backup

To produce VM backups, Veeam Backup & Replication runs backup jobs. A backup job is a collection of settings that define the way backup operations are performed: what data to back up, where to store backups, when to start the backup process, and so on.

One backup job can be used to process multiple VMs, but you can back up each VM with one backup job at a time. If a VM is added to more than one backup job, it will be processed only by the backup job that started earlier.

Creating Backup Jobs

To create a backup job, do the following:

1. Check prerequisites and limitations.
2. Launch the New Backup Job wizard.
3. Specify a job name and description.
4. Select VMs to back up.
5. Specify a backup repository and configure backup settings.
6. Enable guest processing.
7. Create a schedule for the backup job.
8. Finish working with the wizard.

Before You Begin

Before you create a backup job, consider the following limitations:

- You can back up each VM with one backup job at a time. If a VM is already being processed by a backup job, another backup job will not start processing this VM until the currently running backup operation completes.
- You cannot back up a VM being restored. Wait for the restore process to complete, and then start the backup job.
- You cannot back up VMs created from templates as linked clones. Backup of full clones is supported.
- You cannot back up VMs with the same BIOS UUID.
- You cannot include into a backup job a VM that is being backed up by 3rd party software. Wait for the backup process to complete or stop the currently running job manually, and then add the VM to the necessary backup job.
- You cannot include into a backup job a resource pool that does not contain any VMs. Note that after you update a resource pool in the Proxmox VE administration portal, it may take up to 15 minutes for Veeam Plug-in for Proxmox VE to synchronize data between Proxmox VE and Veeam Backup & Replication.
- Veeam Plug-in for Proxmox VE does not support backup of iSCSI disks. If iSCSI disks are attached to a VM included into a backup job, these disks will be skipped from processing.
- Veeam Plug-in for Proxmox VE does not support backup of VM permissions granted to users, user groups and API tokens.
- By default, Veeam Plug-in for Proxmox VE enables deduplication for backed-up data. Due to technical limitations, you cannot disable it while configuring backup jobs.
- By default, backup encryption is disabled for backed-up data. However, you can enable encryption at the repository level as described in the Veeam Backup & Replication User Guide, section Access Permissions.
- Since Veeam Backup & Replication does not allow you to assign information about locations to the Proxmox VE server and workers, job statistics do not include information on the Proxmox VE VM data migration between different geographic regions.

Step 1. Launch New Backup Job Wizard

To launch the **New Job** wizard, do the following:

1. Open the **Home** view.
2. In the inventory pane, select **Jobs**.
3. Right-click the working area and select **Backup** > **Virtual machine** > **Proxmox VE**.

   Alternatively, click **Backup Job** > **Virtual machine** on the ribbon.

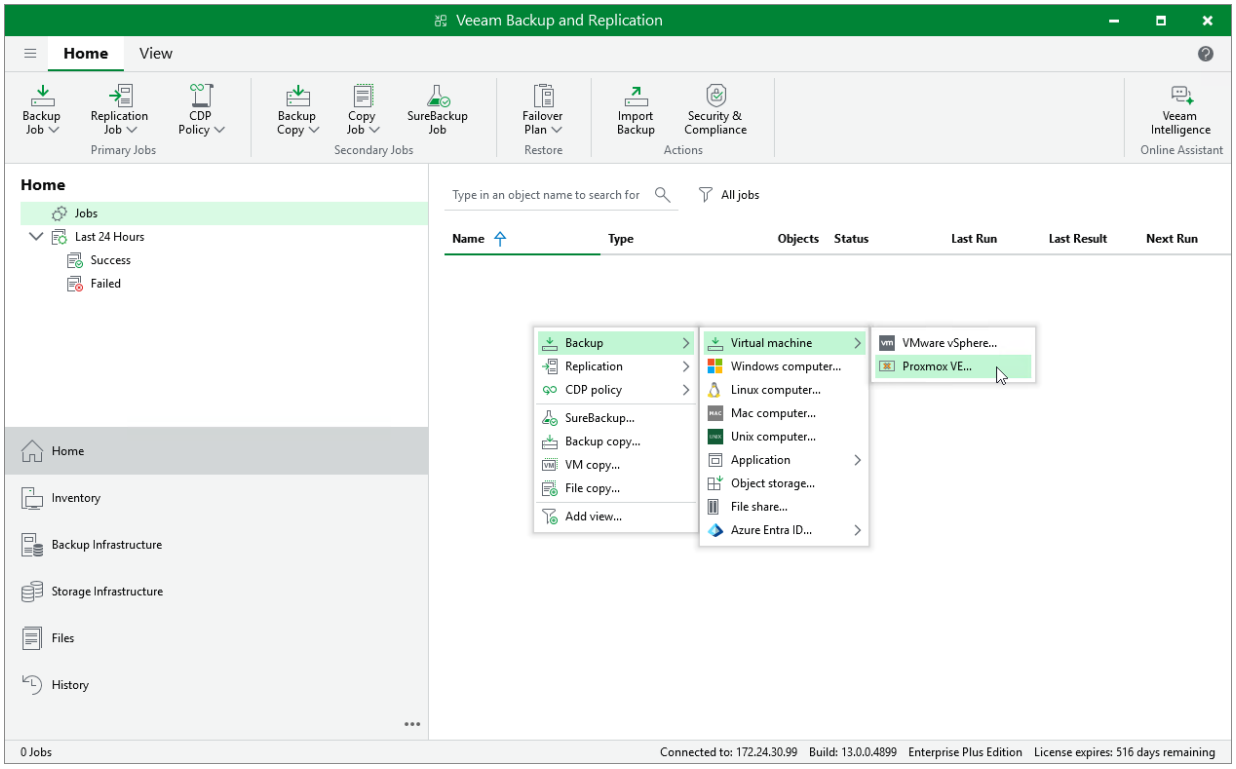

Step 2. Specify Job Name and Description

At the **Name** step of the wizard, use the **Name** and **Description** fields to specify a name for the new backup job and to provide a description for future reference. The job name must be unique in Veeam Backup & Replication.

The maximum length of the name is 40 characters; the following characters are not supported: ~ " # % & * : < > ! ? / \ { | } . ' ` $. The maximum length of the description is 1024 characters.

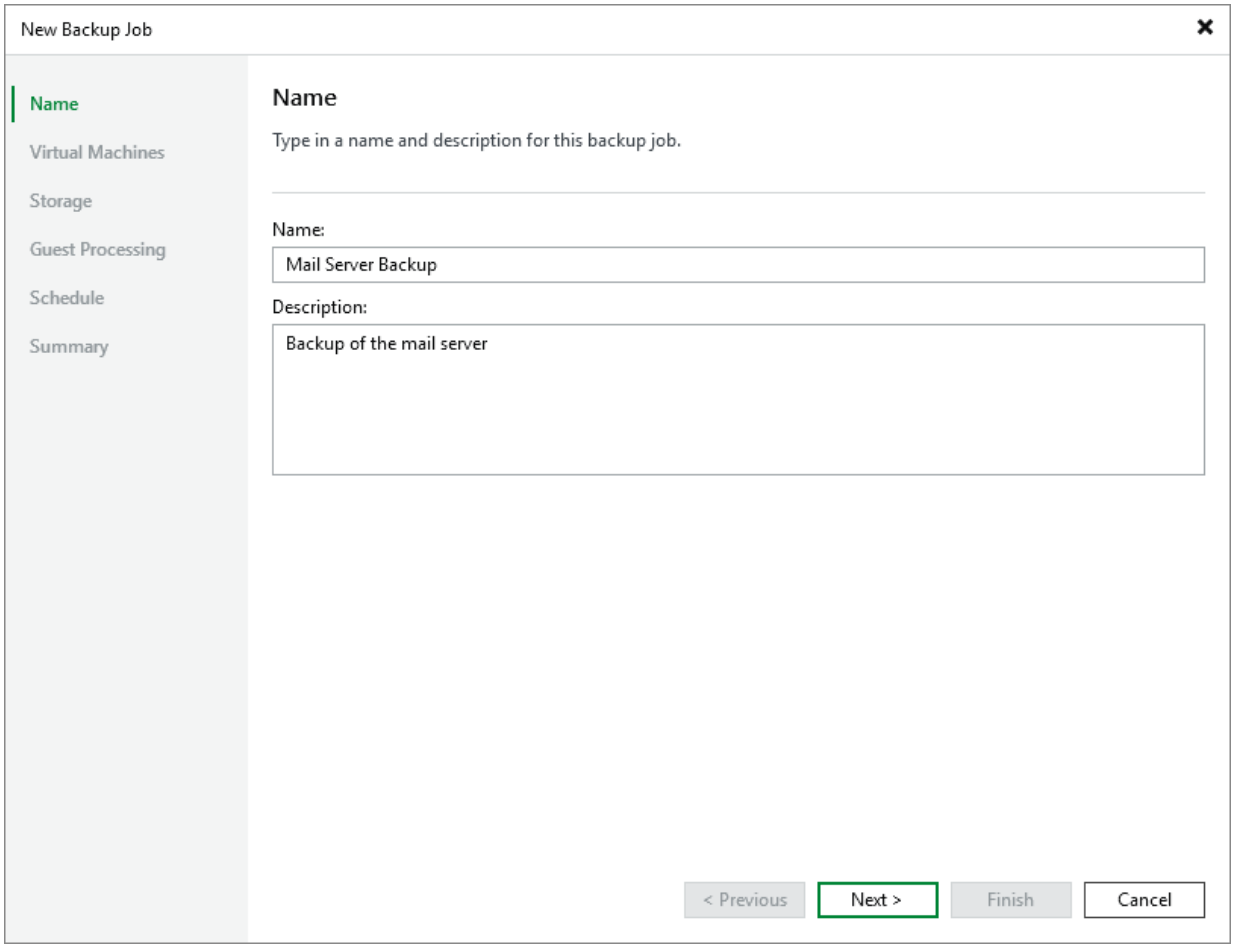

Step 3. Configure Backup Source Settings

At the **Virtual Machines** step of the wizard, specify the backup scope — resources that Veeam Backup & Replication will back up.

Step 3a. Choose Virtual Machines

Specify VMs that will be included into the backup scope:

1. Click **Add**.
2. In the **Add Objects** window, choose whether you want to back up all VMs in the cluster or host, only specific VMs or groups of VMs included into resource pools:

   To view the list of available resource pools, click the **Resource pool** icon on the toolbar at the top right corner of the window. If you add a resource pool to the backup scope, Veeam Backup & Replication will regularly check for new VMs included into the added pool and automatically update the backup job settings to include these VMs in the scope. For a resource pool to be displayed in the list, it must be configured in the Proxmox VE administration portal and must contain at least one VM. For more information on resource pools, see Proxmox VE documentation.

> **TIP**
>
> As an alternative to specifying resources explicitly, you can exclude a number of resources from the backup scope. To do that, click **Exclusions** and specify the VMs that you do not want to back up — the procedure is the same as described for including resources in the backup scope.
>
> Consider that if a resource appears both in the list of included and excluded resources, Veeam Backup & Replication will still not process the resource because the list of excluded resources has a higher priority.

While running the job, Veeam Backup & Replication processes resources in the order they are added to the backup scope. However, you can change the order, for example, if you add some mission-critical VMs to the job and want them to be processed first. To change the processing order, select a resource and use the **Up** or **Down** buttons.

> **NOTE**
>
> Consider the following:
>
> - If you include the same resource into the backup scope multiple times (for example, an individual VM and a resource pool that contains this VM), Veeam Backup & Replication will process this resource only once.
> - If you include a resource pool, node or cluster into the backup scope, VMs in this object are processed at random. To ensure that the VMs are processed in a specific order, you must add them as standalone VMs — not as part of the resource pool, host or cluster.

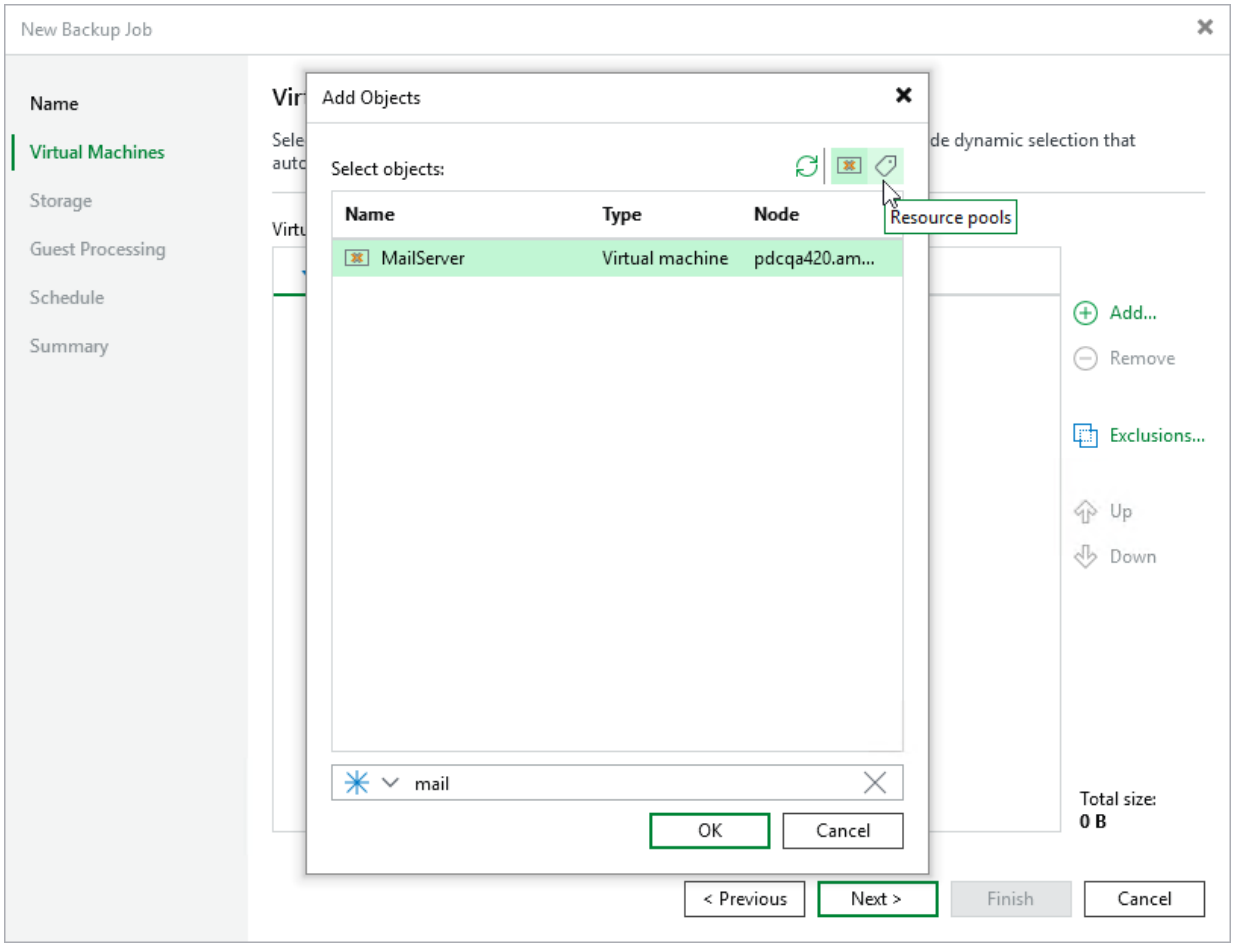

Step 3b. Choose Disks and Volume Groups

By default, jobs process all disks and volume groups attached to VMs included into the backup scope. However, you can instruct Veeam Backup & Replication to back up only specific virtual disks and volume groups related to the selected backup scope:

1. Click **Exclusions**.
2. In the **Exclusions** window, switch to the **Disks** tab and click **Add**.
3. In the **Add Objects** window, select a resource that you have added to the backup scope and click **OK**.
4. Back to the **Exclusions** window, select the resource and click **Edit**.

5. In the **Select Disks** window, select the **Selected Disks** option, click **Add** and choose a bus type of the disks that you want to back up. Then, select the necessary disks.

   Disks that you do not select will be excluded from the backup job.

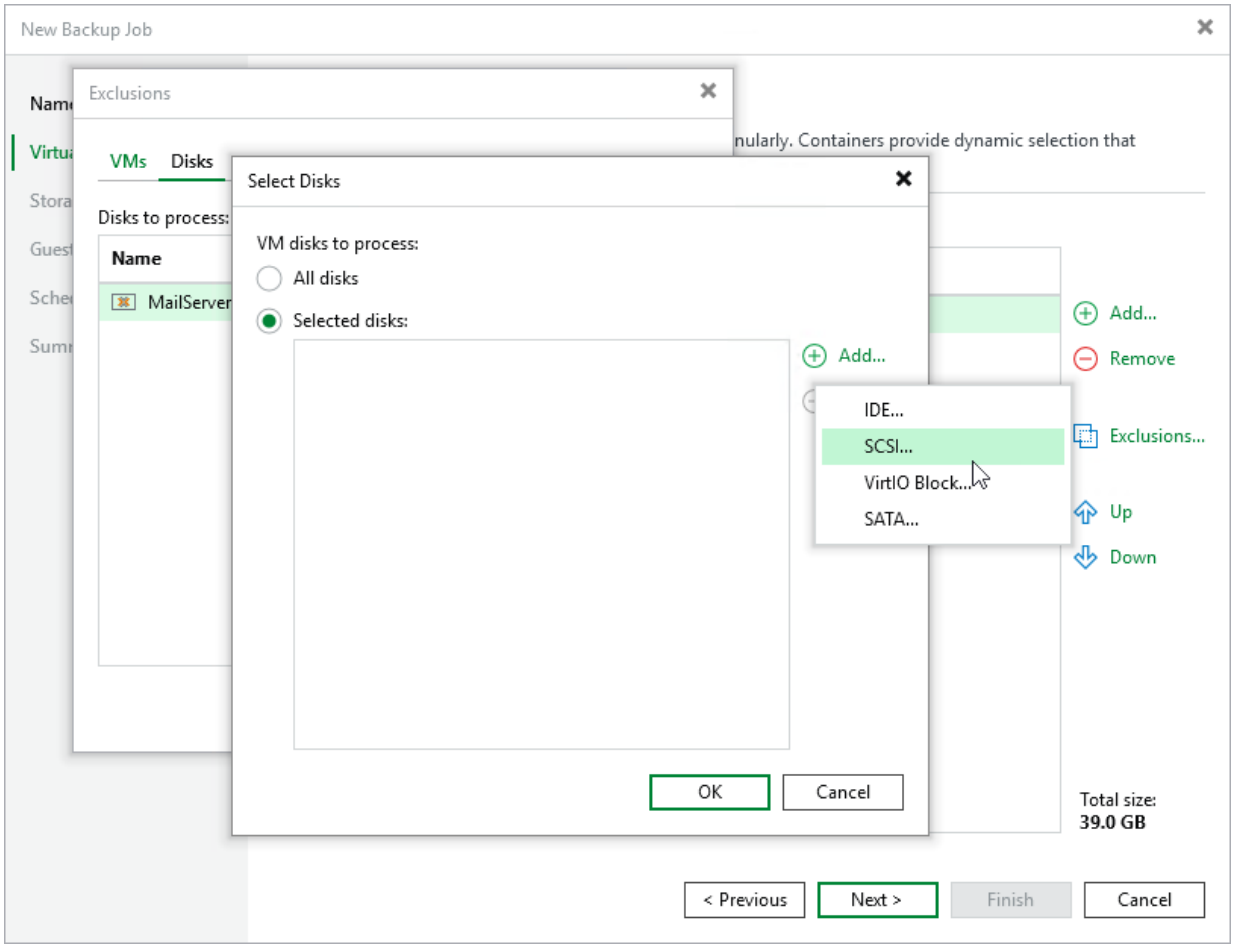

Step 4. Configure Backup Destination Settings

At the **Storage** step of the wizard, do the following:

1. In the **Backup repository** drop-down list, select a backup repository where you want to store backups.

   For a backup repository to be displayed in the list of available repositories, it must be added to the backup infrastructure.

> **NOTE**
>
> Veeam Backup & Replication Community Edition does not support deduplicating storage appliances for storing Proxmox VE VM backups.

2. In the **Retention policy** section, choose how long Veeam Backup & Replication will keep restore points in a backup chain. If a restore point is older than the specified limit, Veeam Backup & Replication will remove it from the chain. For more information on how Veeam Backup & Replication tracks and removes redundant restore points, see Retention Policies.

   Keep in mind that since every backup chain must contain at least 3 restore points, Veeam Backup & Replication may ignore the configured retention policy settings and retain restore points for longer periods of time. For more information, see Backup Retention.

> **NOTE**
>
> If the UUID of a VM changes (for example, if the VM migrates to another cluster), Veeam Backup & Replication will be unable to continue the backup chain for this VM. After you re-add the VM to the backup job, Veeam Backup & Replication will start a new backup chain for it. However, you will still be able to perform restore operations using backups from the old backup chain.

To help you implement a comprehensive backup strategy, Veeam Backup & Replication allows you to enable long-term retention policy for backups and to configure advanced job settings (such as notification settings, health check, active and synthetic full backups).

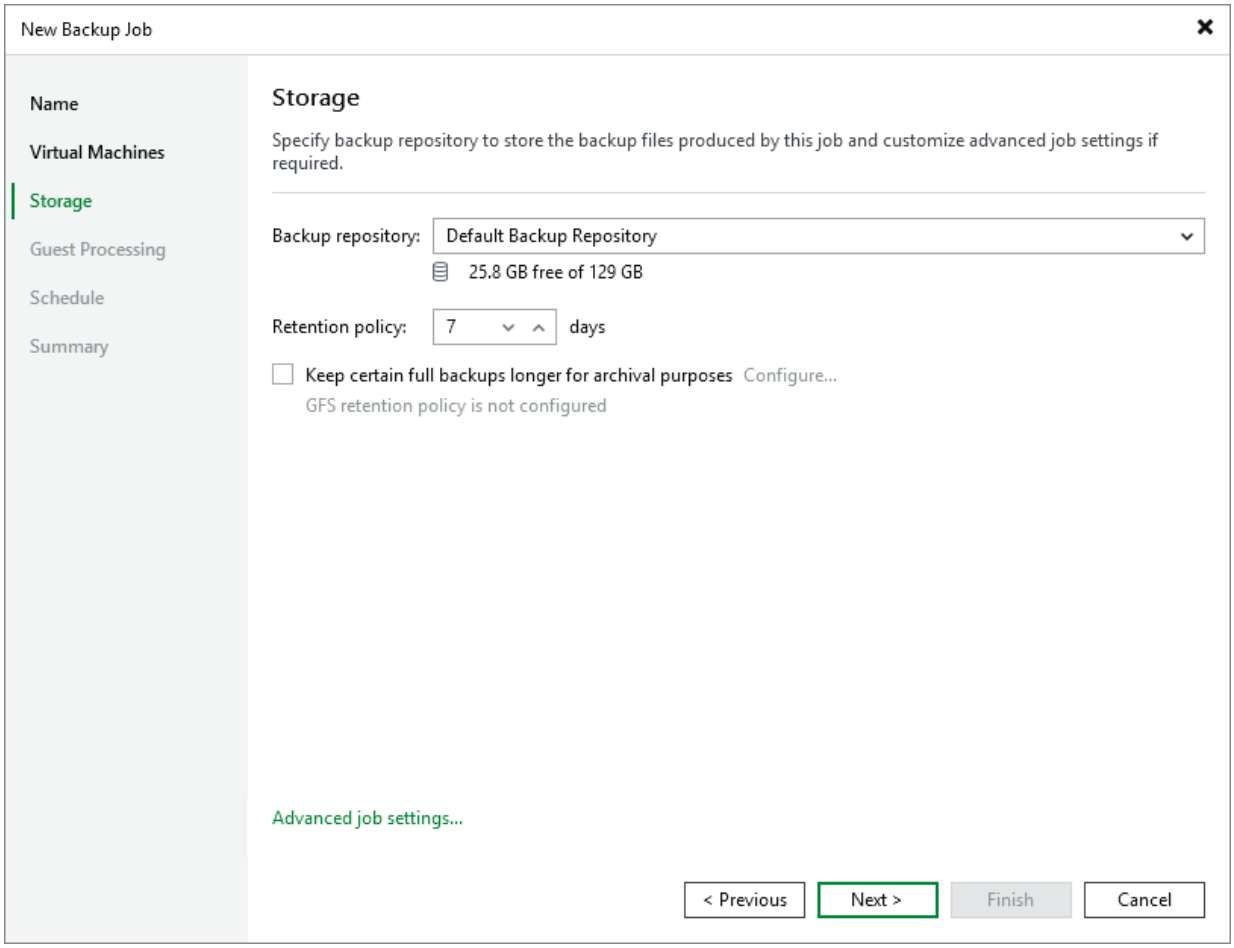

Configuring GFS Policy Schedules

Grandfather-Father-Son (GFS) policy allows you to leverage full backups for long-term retentions instead of creating a new full backup every time. The mechanism simplifies the backup schedule and optimizes the backup performance.

Veeam Backup & Replication re-uses full backups created according to the backup job schedule to achieve the desired retention for a GFS policy schedule (weekly, monthly and yearly). Each full backup is marked with a flag of a specific GFS policy schedule type: the (W) flag is used to mark full backups for the weekly schedule, (M) — monthly, and (Y) — yearly. Veeam Backup & Replication uses these flags to control the retention period for the created full backups. Once a flag of a GFS policy schedule is assigned to a full backup, this full backup can no longer be removed — it is kept for the period defined in the retention settings. When the specified retention period is over, the flag is unassigned from the full backup. If the full backup does not have any other flags assigned, it is removed according to the short-term retention policy settings. For more information on the GFS flag assignment and removal, see the Veeam Backup & Replication User Guide, section Long-Term Retention Policy (GFS).

To configure a GFS policy schedule, select the **Keep certain full backups longer for archival purposes** check box and click **Configure**. Then, specify the following options in the **Configure GFS** window:

- **Keep weekly full backups** — Veeam Backup & Replication will keep a full backup created within a week or on the specific day for a number of weeks.

- **Keep monthly full backups** — Veeam Backup & Replication will keep a full backup created during the specific week for a number of months.
- **Keep yearly full backups** — Veeam Backup & Replication will keep a full backup created in the specific month for a number of years.

After you configure the GFS retention policy settings, schedule active full or synthetic full backups. Otherwise, no new full backups will be automatically produced, and Veeam Backup & Replication will be unable to leverage them for long-term retentions.

> **NOTE**
>
> If you choose an object storage repository to store backups produced by the backup job, you cannot enable synthetic full backups. However, if you configure a GFS policy, synthetic backups will be automatically created according to the specified GFS schedule and marked with an appropriate GFS flag.

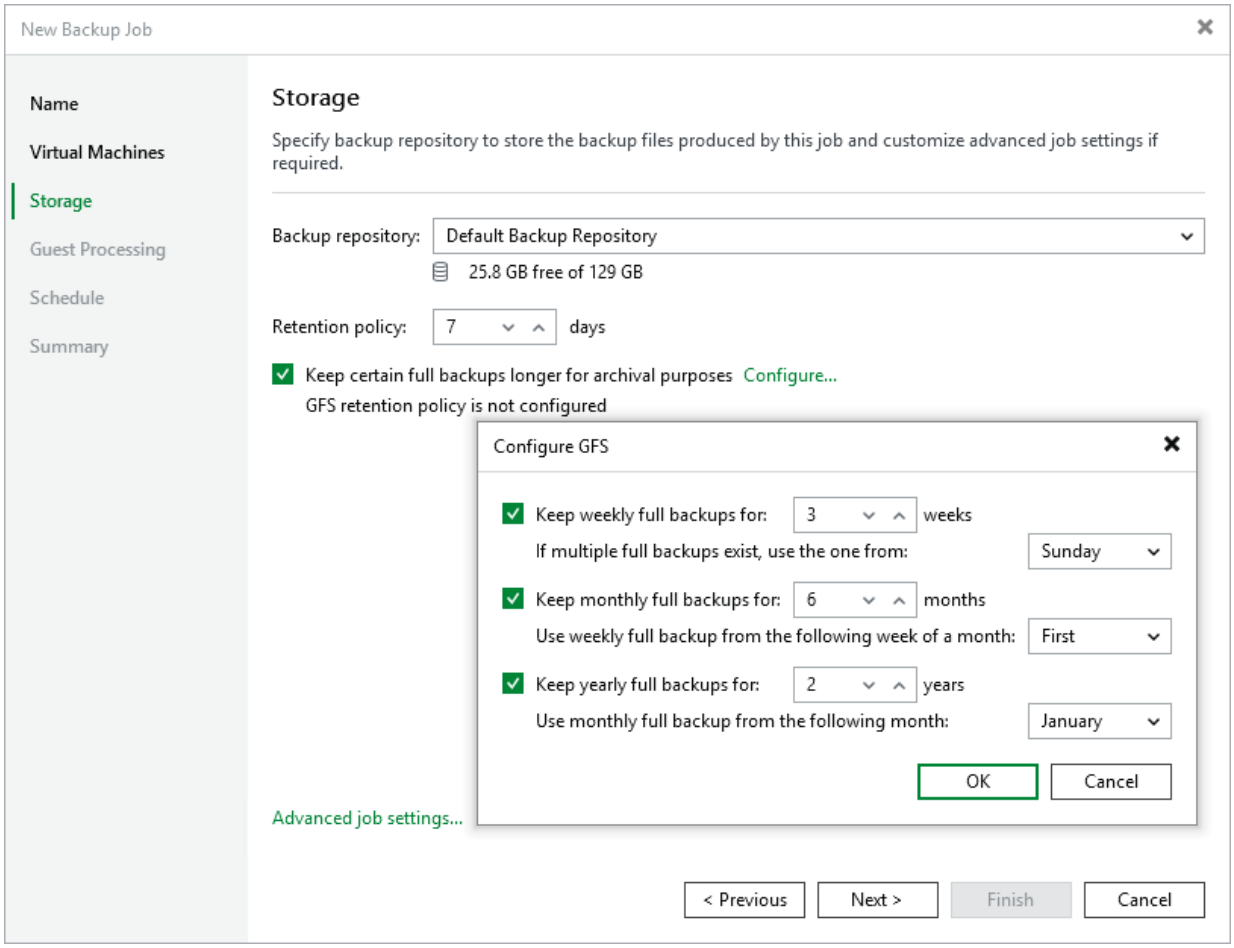

Configuring Advanced Settings

In the **Advanced settings** window, you can schedule full backups, configure health check settings, specify backup file storage settings and enable email notifications.

Backup Settings

To instruct Veeam Backup & Replication to create full backups according to a specific schedule, switch to the **Backup** tab and do the following:

1. To schedule synthetic full backups, select the **Create synthetic full backups periodically** check box, click **Configure** and choose whether you want to create these backups on specific days on a weekly or monthly basis.
2. To schedule active full backups, select the **Create active full backups periodically** check box, click **Configure** and choose whether you want to create these backups on specific days on a weekly or monthly basis.

   Alternatively, you can create active full backups manually when needed. For more information, see Creating Active Full Backups.

IMPORTANT

- Synthetic full backups cannot be scheduled if an object storage repository is selected as the target location for backups.
- Do not schedule synthetic and active full backups to run at the same time. Due to technical limitations, Veeam Backup & Replication will be unable to create synthetic full backups according to the specified schedule.

Health Check Settings

To instruct Veeam Backup & Replication to periodically perform a health check for backups, switch to the **Maintenance** tab, select the **Perform backup files health check (detects and auto-heals corruption)** check box and click **Configure** and specify a schedule for the health check to run.

IMPORTANT

- It is recommended that the backup and health check schedules configured for the job do not overlap to avoid data access issues.
- If you have selected an off-premise cloud object storage repository as the target location for backups at step 4, it is recommended that a helper appliance is configured in the repository settings. Otherwise, additional data transfer costs may occur.

Storage Settings

To specify storage settings for backup files created by the backup job, switch to the **Storage** tab and do the following:

1. To decrease the size of the files, select a compression level from the **Compression level** drop-down list (*None*, *Dedupe-friendly*, *Optimal*, *High* or *Extreme*). For more information on data compression, see the Veeam Backup & Replication User Guide, section Compression and Deduplication.
2. To optimize job performance and storage usage, select a block size from the **Storage optimization** drop-down list. Veeam Backup & Replication will use this size to "split" VM images into separate data blocks when processing VMs — the more data blocks there are, the more time is required to process the VM images. For more information on how data block sizes affect performance, see the Veeam Backup & Replication User Guide, section Storage Optimization.

Notification Settings

To instruct Veeam Backup & Replication to send email notifications on the backup job results, switch to the **Notifications** tab, select the **Send email notifications** check box and specify an email address of a recipient; use a semicolon to separate multiple recipient addresses. For Veeam Backup & Replication to be able to send email notifications, you must configure a mail server as described in section Configuring Email Settings.

> **NOTE**
>
> Email notifications on the backup job results will be also sent to recipients configured in the global notification settings.

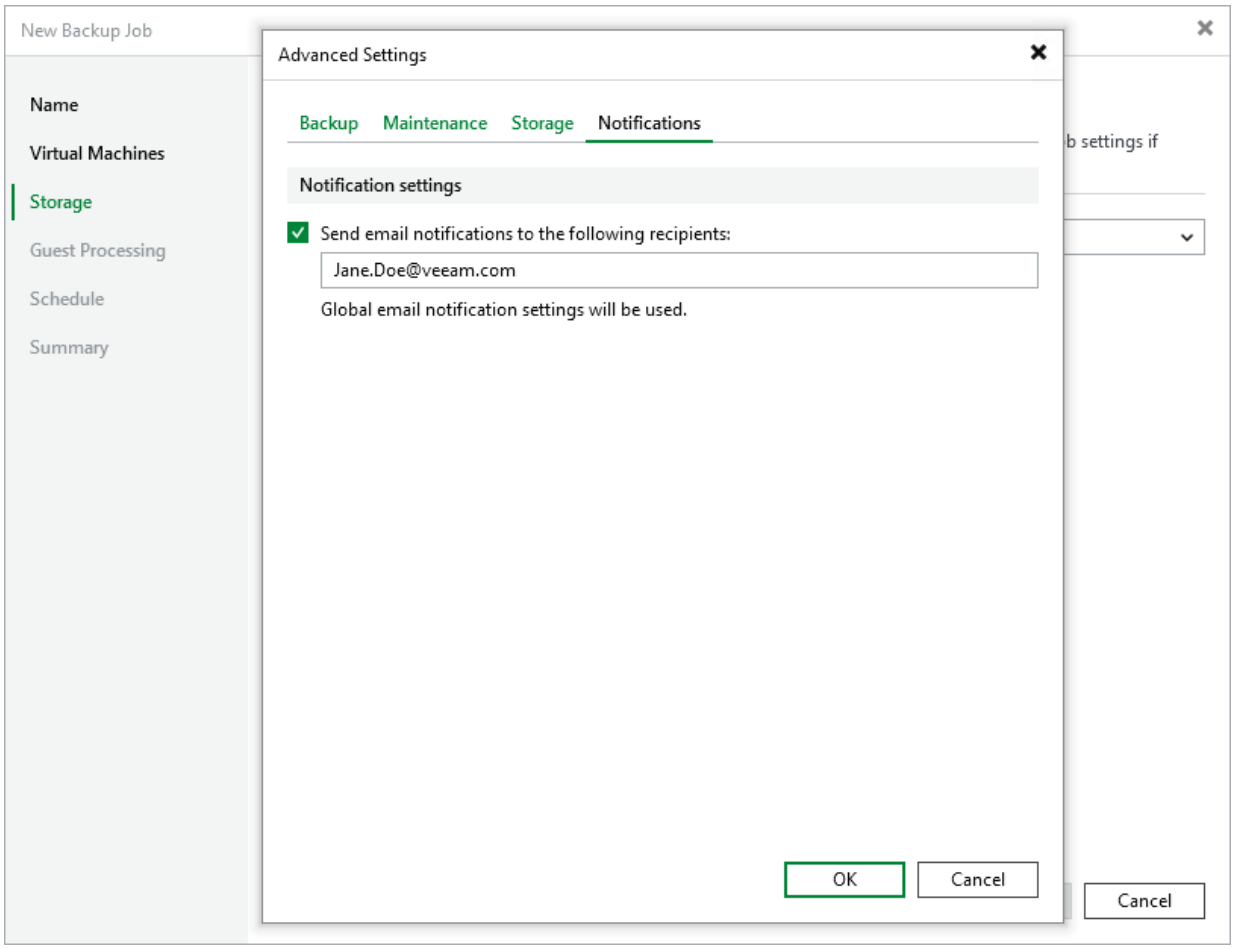

How Health Check Works

When Veeam Backup & Replication saves a new backup restore point to a backup repository, it calculates CRC values for metadata in the backup chain and saves these values to the chain metadata, together with the instance data. When performing a health check, Veeam Backup & Replication verifies the availability of data blocks and uses the saved values to ensure that the restore points being verified are consistent.

On the day scheduled for a health check to run, Veeam Backup & Replication starts a new health check session. For each restore point in the standard backup chain, Veeam Backup & Replication calculates CRC values for backup metadata and compares them to the CRC values that were previously saved to the restore point. Veeam Backup & Replication also checks whether data blocks that are required to rebuild the restore point are available.

If Veeam Backup & Replication does not detect data inconsistency, the health check session completes successfully. Otherwise, the session completes with an error. Depending on the detected data inconsistency, Veeam Backup & Replication performs the following operations:

- If the health check detects corrupted metadata in a full or incremental restore point, Veeam Backup & Replication marks the backup chain as corrupted in the configuration database. During the next backup job session, Veeam Backup & Replication copies the full instance image, creates a full restore point in the backup repository and starts a new backup chain in the backup repository.
- If the health check detects corrupted disk blocks in a full or an incremental restore point, Veeam Backup & Replication marks the restore point that includes the corrupted data blocks and all subsequent incremental restore points as incomplete in the configuration database. During the next backup job session, Veeam Backup & Replication copies not only those data blocks that have changed since the previous backup session but also data blocks that have been corrupted, and saves these data blocks to the latest restore point that has been created during the current session.

Step 5. Specify Guest Processing Options

At the **Guest Processing** step of the wizard, you can specify the following settings:

- Enable application-aware processing — to create transactionally consistent backups that will guarantee proper recovery of VM applications, without data loss.

  For VMs running Microsoft SQL Server, Oracle Server or PostgreSQL Server applications, you can also instruct Veeam Backup & Replication to periodically back up transaction logs. This will allow you to restore your databases to specific points in time as described in the Veeam Enterprise Manager User Guide, section Restoring Point-in-Time State.

- Enable guest file system indexing and malware detection — to create a catalog of guest OS files that will allow you to search for specific items during file-level restore. This will also allow you to receive reports about malware files and suspicious system activity detected on VMs included into the backup scope.

- Choose guest interaction proxies — to select specific servers that Veeam Backup & Replication will use when communicating with guest OSes of VMs included into the backup scope.

- Manage VM guest OS credentials — to specify credentials that Veeam Backup & Replication will use to access guest OSes of all VMs included into the backup scope.

Considerations and Limitations

If you enable application-aware processing or guest files system indexing, consider the following:

- Veeam Plug-in for Proxmox VE will not be able to use Kerberos authentication while connecting to guest OSes of the processed VMs.

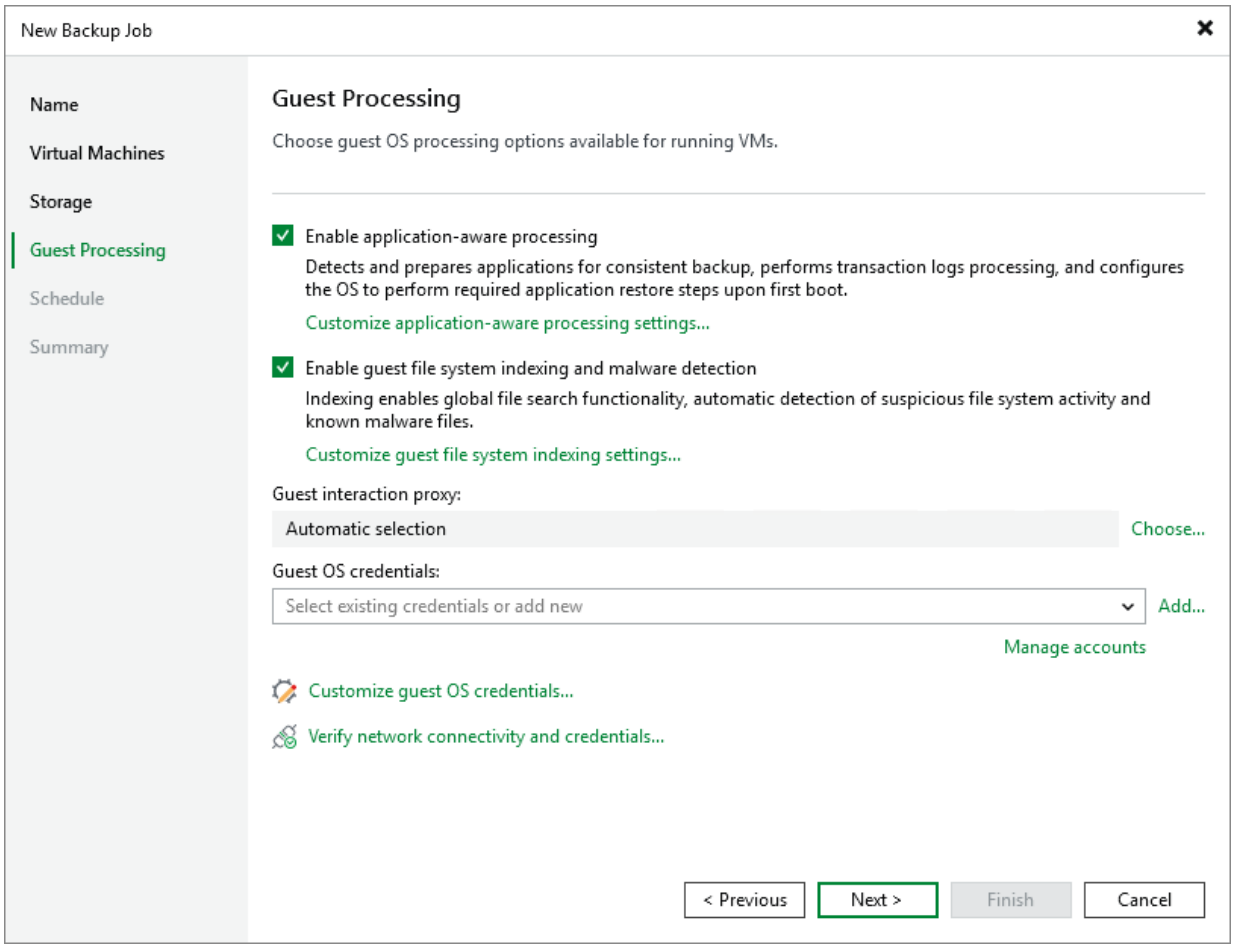

Related Topics

- Requirements and limitations for PostgreSQL WAL files backup
- Requirements and limitations for Oracle archived redo logs backup

Step 5a. Enable Application-Aware Processing

To restore your applications without data loss, you must allow Veeam Backup & Replication to create application-consistent backups. To do that, select the **Enable application-aware processing** check box at the **Guest Processing** step of the wizard.

When creating application-consistent backups, Veeam Backup & Replication takes transactionally consistent VM snapshots while no write operations occur on VM disks. To do that, Veeam Backup & Replication quiesces applications on the processed VMs and creates a consistent view of application data:

- To quiesce VSS-aware applications running on Windows-based VMs (such as MS SQL, MS Exchange, Microsoft Active Directory and Microsoft SharePoint), Veeam Backup & Replication leverages the Microsoft VSS technology.
- To quiesce applications running on Linux-based VMs and non-VSS-aware applications running on Windows-based VMs, Veeam Backup & Replication runs custom scripts before and after the snapshot creation.

For more information on supported applications that can be protected with application-consistent backups, see the Veeam Backup & Replication User Guide, section Supported Platforms and Applications.

Processing Transaction Logs

If you enable application-aware processing, Veeam Backup & Replication will back up and truncate transaction logs produced by VM applications every time the backup job starts. To change this behavior, you can do either of the following:

- Instruct Veeam Backup & Replication not to process and truncate logs. This will allow third-party backup solutions to perform VM guest-level backup and to maintain consistency of the database state.
- Instruct Veeam Backup & Replication to back up and truncate transaction logs more often. This will allow you to use application-consistent backups to restore your MS SQL, Oracle and PostgreSQL databases to specific points in time.

To configure log processing settings, complete the following steps:

1. Click **Customize application-aware processing settings**.
2. In the **Application-Aware Processing Options** window, select the necessary resource and click **Edit**. You can configure guest processing settings for multiple resources at a time.

   If you want to configure processing settings for a specific VM that is included into a resource pool, host or cluster, you must configure those settings separately. To do that, click **Add**, choose the necessary VM and click **Edit**.
3. In the **Processing Settings** window, do the following:
   - To specify how Veeam Backup & Replication will process transaction logs of VSS-aware applications, select the **Process transaction logs with this job** option on the **General** tab, switch to the **SQL** tab and follow the instructions provided in section Specifying Microsoft SQL Server Transaction Log Settings.

     If you do not want Veeam Backup & Replication to process and truncate transaction logs of VSS-aware applications, select the **Perform copy only** option. However, with this option selected, the backup job will produce copy-only backups that cannot be used to restore MS SQL databases to specific points in time. For more information on copy-only backups, see Microsoft Docs.
   - To specify how Veeam Backup & Replication will process transaction logs of Oracle Server applications, switch to the **Oracle** tab and follow the instructions provided in section Specifying Oracle Archived Redo Log Settings.

- To specify how Veeam Backup & Replication will process transaction logs of PostgreSQL Server applications, switch to the **PostgreSQL** tab and follow the instructions provided in section Specifying PostgreSQL WAL Files Settings.
- To specify scripts that Veeam Backup & Replication will use to quiesce non-VSS-aware applications, switch to the **Scripts** tab and follow the instructions provided in section Specifying Pre-Freeze and Post-Thaw Scripts.

> **TIP**
>
> To instruct Veeam Backup & Replication not to perform application-aware processing for the selected resource at all, select the **Disable application processing** option.

Handling Application-Aware Processing Errors

By default, Veeam Backup & Replication requires application-aware processing to finish without errors for the backup job to complete successfully. In case of an error, Veeam Backup & Replication terminates the backup operation, and the backup job will not process transaction logs until a new image-level backup is created for each of the VMs included into the backup scope.

To change this behavior and instruct Veeam Backup & Replication to proceed with the backup operation, creating a crash-consistent backup instead of an application-consistent backup, switch to the **General** tab of the **Processing Settings** window and select the **Try application processing, but ignore failures option**.

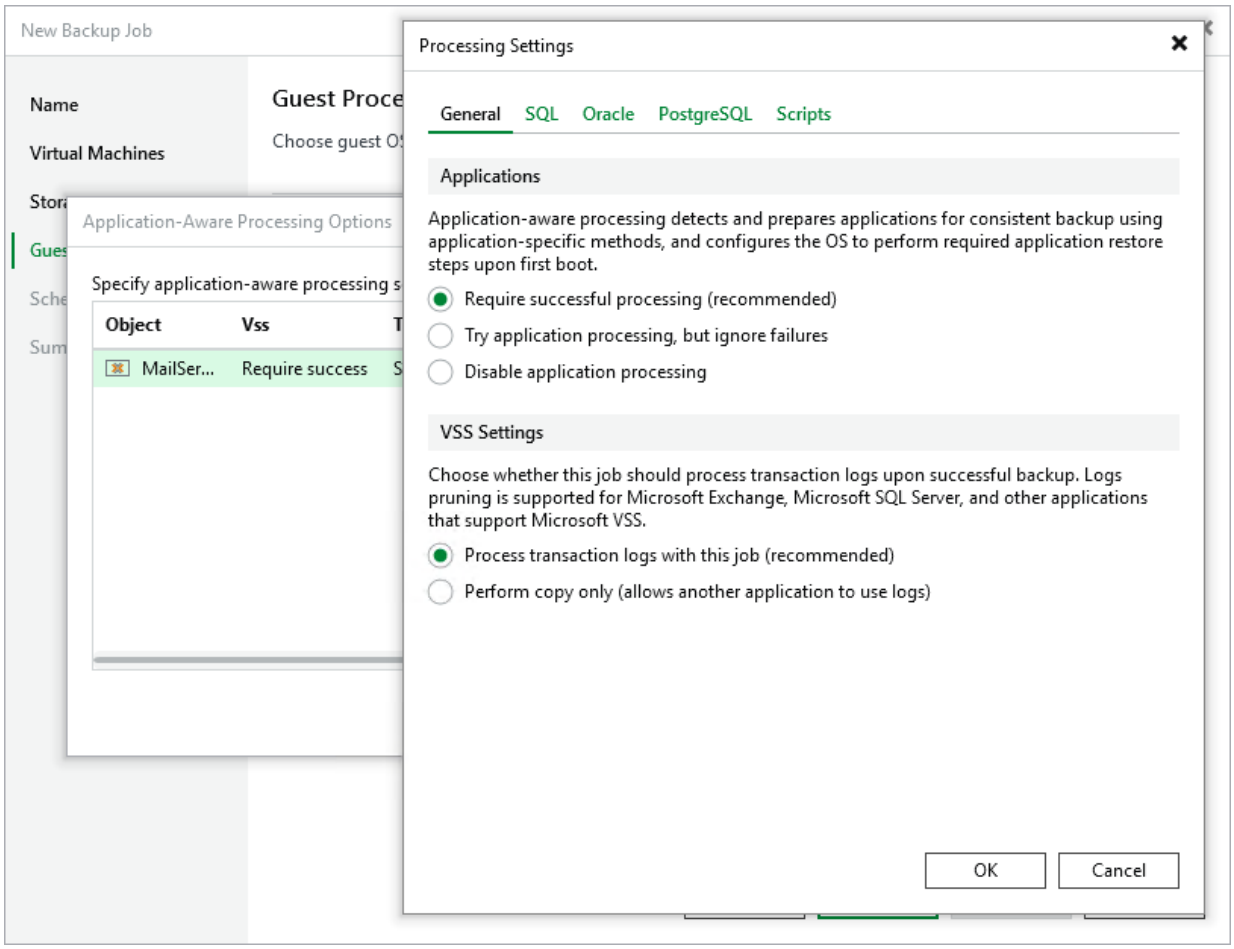

Specifying Microsoft SQL Server Transaction Log Settings

By default, Veeam Backup & Replication creates application-consistent image-level backups of VMs running the Microsoft SQL Server application and truncates transaction logs after each successfully completed backup session — this will allows you to restore Microsoft SQL Server databases using specific backups. To protect mission-critical Microsoft SQL Server databases, you can instruct Veeam Backup & Replication to create secondary restore points with transaction logs in addition to primary image-level backups — this will allow you to restore your databases to specific points in time.

NOTES

- Veeam Backup & Replication stores image-level backups and transaction log backups in the same repository.
- If Veeam Backup & Replication fails to produce a primary image-level backup, no secondary transaction log backups will be created.

To back up Microsoft SQL Server transaction logs periodically, do the following:

1. Switch to the **SQL** tab and select the **Backup logs periodically** option.
2. In the **Backup logs every** field, specify how frequently you want transaction logs to be backed up. The maximum field value is 480 minutes.
3. In the **Retain log backups section**, choose either of the following options:
   - Select the **Until the corresponding image-level backup is deleted** option if you want to remove transaction log backups and the related image-level backups at the same time, according to the retention policy settings specified at step 4.
   - Select the **Keep only last <N> days of log backups** option if you want to retain transaction logs for a specific time period, regardless of the retention policy settings specified for image-level backups. Note that image-level backups must always be kept for a longer period than the related transaction log backups.

   For more information on how Veeam Backup & Replication retains transaction logs, see the Veeam Backup & Replication User Guide, section Microsoft SQL Server Log Backup.
4. In the **Log shipping servers** section, choose whether you want to use a specific Windows server to transfer transaction log backups or let Veeam Backup & Replication choose it automatically to reduce the load on the backup server.

   By default, Veeam Backup & Replication automatically chooses a log shipping server for each of the processed VMs based on network settings and rules listed in the Veeam Backup & Replication User Guide, section Log Shipping Servers. You can also manually limit the list of machines that may be used as log shipping servers — to do that, click **Choose**, select the **Use the specified servers only** option and then select check boxes next to the necessary Windows servers.

   For a Windows server to be displayed in the list of available log shipping servers, it must be added to the backup infrastructure as described in the Veeam Backup & Replication User Guide, section Adding Microsoft Windows Servers. Keep in mind that the list will also include Linux servers added to the backup infrastructure; however, Linux servers cannot be used as log shipping servers due to technical limitations in the current version.

TIPS

- It is recommended that you choose at least 2 log shipping servers for load balancing and high availability purposes.
- It is recommended that you do not choose servers that are engaged in permanent tasks consuming resources (such as WAN accelerators or backup servers).

You can also choose not to truncate logs at all. However, keep in mind that this option requires databases to use the simple recovery model. Otherwise, transaction logs may grow large and increase the storage space consumption significantly. For more information on recovery models used by Microsoft SQL databases, see Microsoft Docs.

Considerations and Limitations

When you configure transaction log settings, consider the following:

- If a processed VM runs Microsoft SQL Server along with Oracle Server and transaction log backup is enabled for both applications, the Microsoft SQL Server transaction logs will not be backed up — Veeam Backup & Replication will create transaction log backups for the Oracle Server application only.
- If a processed VM runs Microsoft SQL Server that hosts the Veeam Backup & Replication configuration database, its transaction logs will not be backed up:
  - If the Microsoft SQL Server has the SQL Server Always On availability groups feature disabled, the configuration database will be excluded from application-aware processing automatically.
  - If the Microsoft SQL Server has the SQL Server Always On availability groups feature enabled, you will have to exclude the configuration database from application-aware processing manually, as described in this Veeam KB article.

  For more information on the SQL Server Always On availability group feature, see Microsoft Docs.

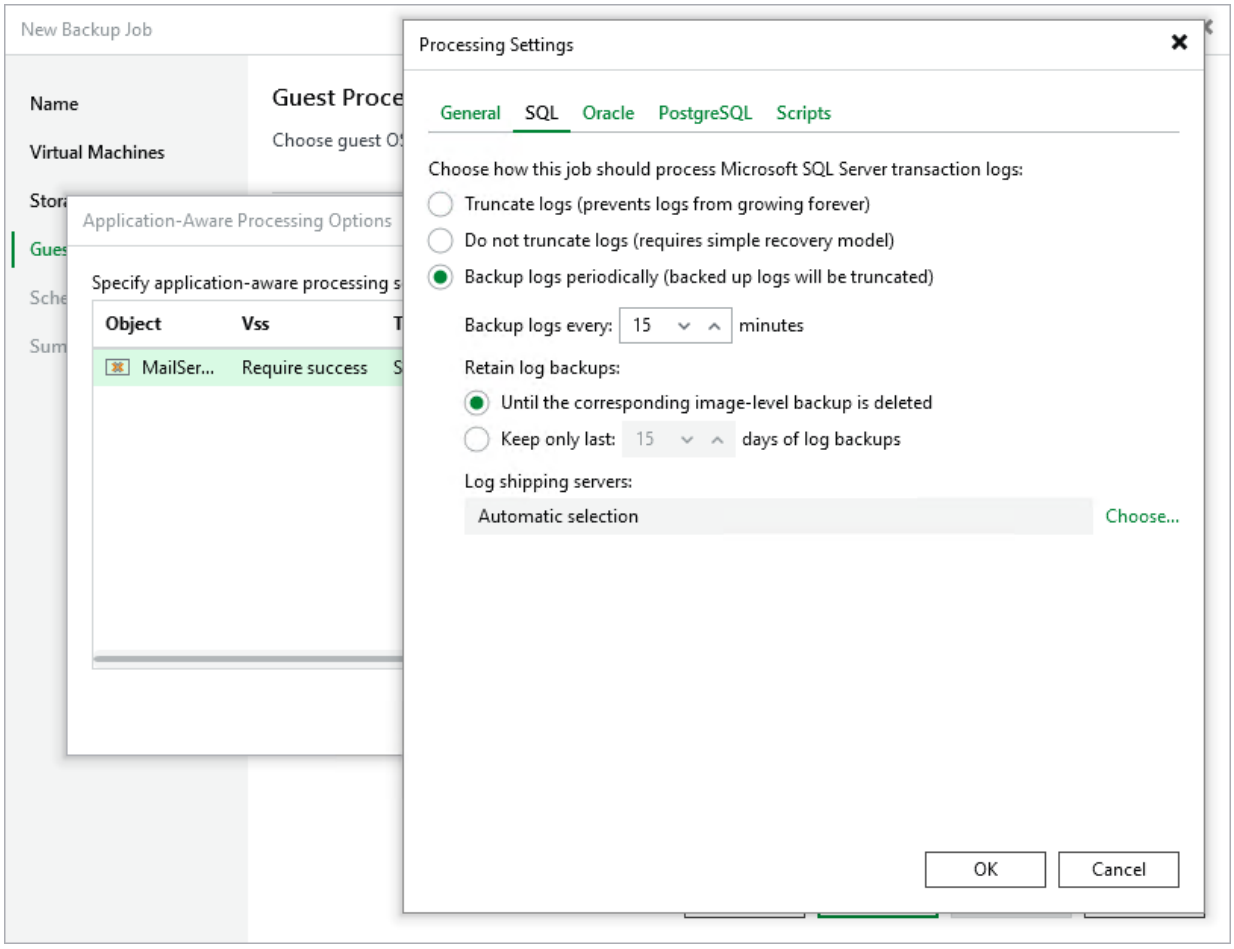

Specifying Oracle Archived Redo Log Settings

By default, Veeam Backup & Replication creates application-consistent image-level backups of VMs running the Oracle application and does not truncate archived redo logs after each successfully completed backup session — this allows you to restore Oracle databases using specific backups. To protect mission-critical Oracle databases, you can instruct Veeam Backup & Replication to create secondary restore points with archived redo logs in addition to primary image-level backups — this will allow you to restore your databases to specific points in time.

NOTES

- Veeam Backup & Replication stores image-level backups and archived redo log backups in the same repository.
- If Veeam Backup & Replication fails to produce a primary image-level backup, no secondary archived redo log backups will be created.

To back up Oracle archived redo logs periodically, do the following:

1. Switch to the **Oracle** tab.
2. In the **Backup logs every** field, specify how frequently you want archived redo logs to be backed up. The maximum field value is 480 minutes.
3. In the **Retain log backups** section, choose either of the following options:
   - Select the **Until the corresponding image-level backup is deleted** option if you want to remove archived redo log backups and the related image-level backups at the same time, according to the retention policy settings specified at step 4.
   - Select the **Keep only last <N> days of log backups** if you want to retain archived redo log backups for a specific time period, regardless of the retention policy settings specified for image-level backups. Note that archived redo logs backups must always be retained for a longer period than image-level backups.

   For more information on how Veeam Backup & Replication retains archived redo logs, see the Veeam Backup & Replication User Guide, section Retention for Archived Log Backup.
4. In the **Log shipping servers** section, decide whether you want to use a specific server to transfer archived redo logs backups or let Veeam Backup & Replication choose it automatically to reduce the load on the backup server.

   By default, Veeam Backup & Replication automatically chooses a log shipping server for each of the processed VMs based on network settings and rules listed in the Veeam Backup & Replication User Guide, section Log Shipping Servers. You can also manually limit the list of machines that may be used as log shipping servers — to do that, click **Choose**, select the **Use the specified servers only** option and then select check boxes next to the necessary servers.

   For a server to be displayed in the list of available log shipping servers, it must be added to the backup infrastructure as described in the Veeam Backup & Replication User Guide, sections Adding Microsoft Windows Servers and Adding Linux Servers. Keep in mind that the list will also include Linux servers added to the backup infrastructure; however, Linux servers cannot be used as log shipping servers for processing Windows-based VMs due to technical limitations in the current version.

TIPS

- It is recommended that you choose at least 2 log shipping servers for load balancing and high availability purposes.
- It is recommended that you do not choose servers that are engaged in permanent tasks consuming resources (such as WAN accelerators or backup servers).

You can choose to keep the default **Do not delete archived** logs option, but in this case archived redo logs may grow large and increase the storage space consumption significantly. That is why it is recommended that you choose to remove archived redo logs that are older than a specific time limit or whose size exceeds a specific storage threshold. Keep in mind that the selected option will apply to logs of each processed Oracle database individually — and only after the backup job completes successfully.

Configuring Access to Oracle Data

To access databases of the processed Oracle applications, Veeam Backup & Replication uses accounts with SYSDBA privileges — by default, these are the accounts you specify for accessing the VM guest OSes. To change this behavior, you can choose another account from the **Specify Oracle account with SYSDBA privileges** drop-down list.

For an account to be displayed in the list of available accounts, it must be added to the Credentials Manager as described in the Veeam Backup & Replication User Guide, section Credentials Manager. If you have not added the necessary account to the Credentials Manager beforehand, you can do it without closing the **New Job** wizard. To do that, click either the **Manage accounts** link or the **Add** button, and specify the user name, password and description in the **Credentials** window.

Considerations and Limitations

When you configure transaction log settings, consider the following:

- For Oracle databases running in the NOARCHIVELOG mode, Veeam Backup & Replication is not able to create restore points with archived redo logs — only image-level backups will be created. For more information on how to choose between database modes, see Oracle documentation.
- For Veeam Backup & Replication to be able to access Oracle command-line tools when performing application-aware processing for Windows-based VMs, the *%ORACLE_HOME%\bin* directory must be added to the `PATH` system variable. For more information on how to set Oracle environment variables, see Oracle documentation.

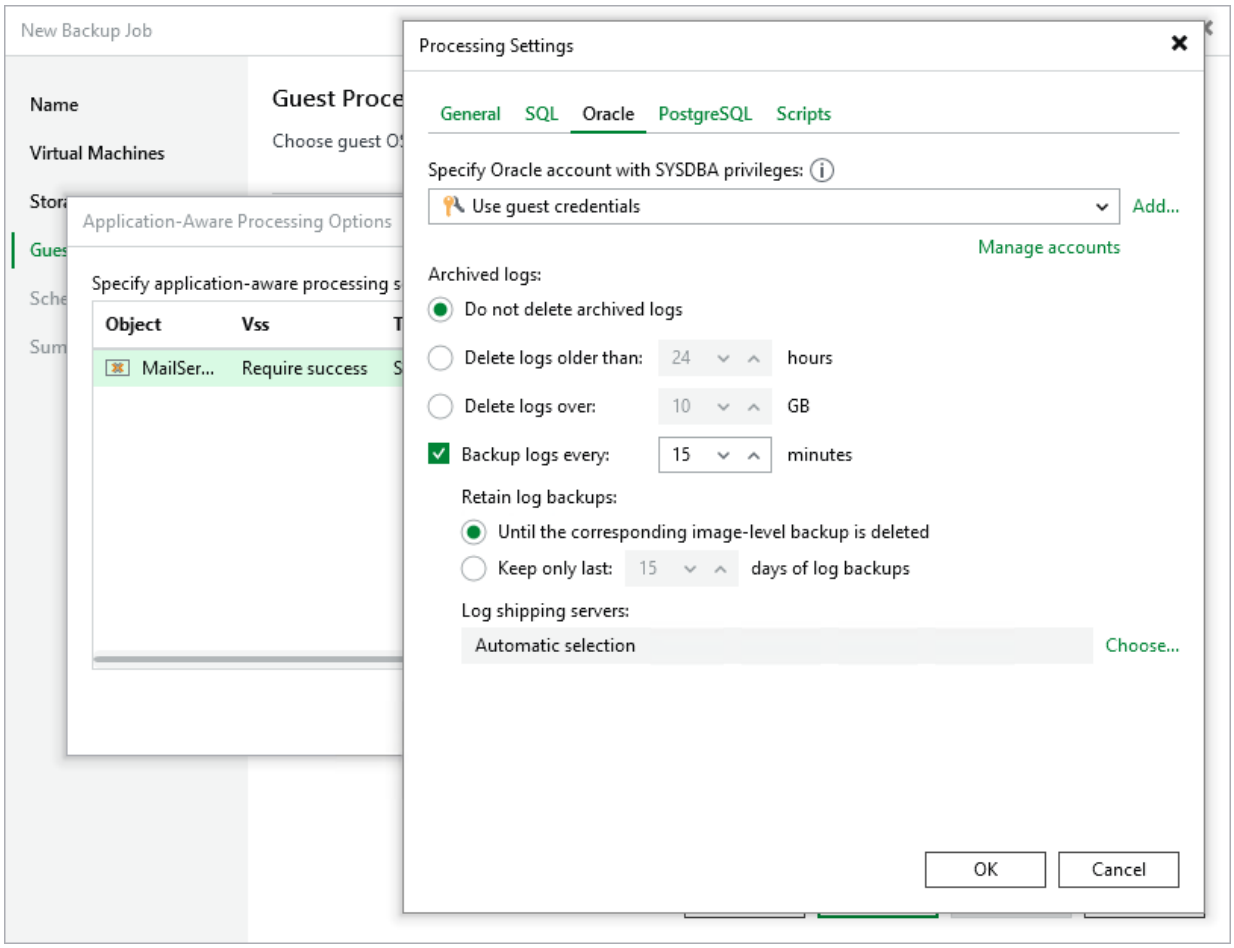

Specifying PostgreSQL WAL Files Settings

By default, Veeam Backup & Replication creates application-consistent image-level backups of VMs running the PostgreSQL Server application and does not truncate write ahead logs (WAL) after each successfully completed backup session — this allows you to restore PostgreSQL Server databases using specific backups. To protect mission-critical PostgreSQL Server databases, you can instruct Veeam Backup & Replication to create secondary restore points with WAL logs in addition to primary image-level backups — this will allow you to restore your databases to specific points in time.

NOTES

- Veeam Backup & Replication stores image-level backups and WAL log backups in the same repository.
- If Veeam Backup & Replication fails to produce a primary image-level backup, no secondary WAL log backups will be created.

To back up PostgreSQL WAL logs periodically, do the following:

1. Switch to the **PostgreSQL** tab.
2. In the **Backup logs every** field, specify how frequently you want WAL logs to be backed up. The maximum field value is 480 minutes.
3. In the **Retain log backups** section, choose either of the following options:
   - Select the **Until the corresponding image-level backup is deleted** option if you want to remove WAL log backups and the related image-level backups at the same time, according to the retention policy settings specified at step 4 of the wizard.

- Select the **Keep only last <N> days of log backups** if you want to retain WAL log backups for a specific time period, regardless of the retention policy settings specified for image-level backups. Note that WAL log backups must always be retained for a longer period than image-level backups.

For more information on how Veeam Backup & Replication retains WAL logs, see the Veeam Backup & Replication User Guide, section Retention for PostgreSQL WAL Files.

4. In the **Temporary location for archive logs** section, specify the path to a folder on the PostgreSQL machine where Veeam Backup & Replication will temporarily store archive logs until they are backed up.

   Keep in mind that you must create the folder beforehand manually. Also, make sure that there is enough free space in this folder for the log files and required permissions are granted to the user account.

5. In the **Log shipping servers** section, decide whether you want to use a specific server to transfer WAL log backups or let Veeam Backup & Replication choose it automatically to reduce the load on the backup server.

   By default, Veeam Backup & Replication automatically chooses a log shipping server for each of the processed VMs based on network settings and rules listed in the Veeam Backup & Replication User Guide, section Log Shipping Servers. You can also manually limit the list of machines that may be used as log shipping servers — to do that, click **Choose**, select the **Use the specified servers only** option and then select check boxes next to the necessary servers.

   For a server to be displayed in the list of available log shipping servers, it must be added to the backup infrastructure as described in the Veeam Backup & Replication User Guide, sections Adding Microsoft Windows Servers and Adding Linux Servers.

> **TIPS**
> - It is recommended that you choose at least 2 log shipping servers for load balancing and high availability purposes.
> - It is recommended that you do not choose servers that are engaged in permanent tasks consuming resources (such as WAN accelerators and backup servers).

Configuring Access to PostgreSQL Data

To access databases of the processed PostgreSQL Server instances, Veeam Backup & Replication uses accounts with superuser privileges — by default, these are the accounts you specify for accessing the VM guest OSes. To change this behavior, you can choose another account from the **Specify PostgreSQL account with superuser privileges** drop-down list. For an account to be displayed in the list of available accounts, it must be added to the Credentials Manager as described in the Veeam Backup & Replication User Guide, section Credentials Manager. If you have not added the necessary account to the Credentials Manager beforehand, you can do it without closing the **New Job** wizard. To do that, click either the **Manage accounts** link or the **Add** button, and specify the user name, password and description in the **Credentials** window.

After you choose an account, you must also explicitly specify whether this account is a PostgreSQL database account (whose password is stored either in the Credentials Manager or in a configuration file) or a system account. In the latter case, make sure that the account has all the permissions required to access the PostgreSQL Server instance; for more information on the required permissions, see Planning and Preparation.

IMPORTANT

Consider the following:

- If the password that is stored in the Credential manager is empty, Veeam Backup & Replication will be able to use this account by leveraging username map authentication. In this case, the `SYSTEM-USERNAME` variable value will be set to the name of the account used to access the VM guest OS, while the `PG-USERNAME` variable value will be set to the name of the account that you have selected from the **Specify PostgreSQL account with superuser privileges** drop-down list. For more information on username map authentication, see PostgreSQL documentation.
- If the password is stored in a .PGPASS configuration file, make sure that it is located in the */home* directory of the selected account. For more information on .PGPASS files, see PostgreSQL documentation.

Depending on the scope of resources that you have specified at step 5a of the wizard, Veeam Backup & Replication will use the selected account in the following way:

- If the scope includes an individual VM, the account will be used to access the PostgreSQL Server instance running on this specific VM.
- If the scope includes multiple individual VMs, the account will be used to access the PostgreSQL Server instance running on each of these VMs.
- If the scope includes a VM container (such as resource pool, host or cluster), the account will be used to access every PostgreSQL Server instance running on VMs in this container.

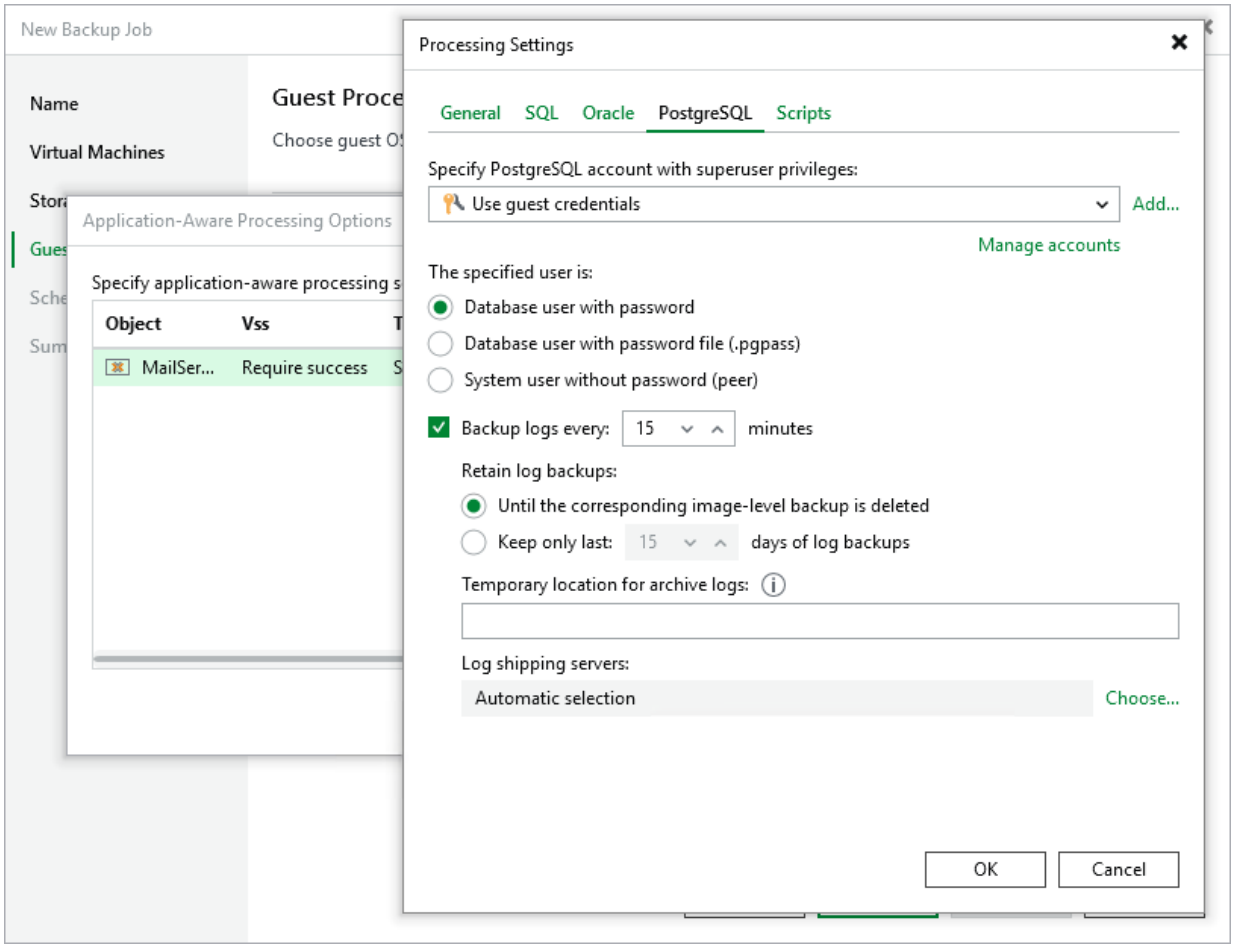

Specifying Pre-Freeze and Post-Thaw Scripts

If you plan to back up VMs running applications that do not support VSS, you can specify what scripts Veeam Backup & Replication must use to quiesce the VM. The pre-freeze script quiesces the VM file system and application data to bring the VM to a consistent state before Veeam Backup & Replication triggers a VM snapshot. After the VM snapshot is created, the post-thaw script brings the VM and applications to their initial state.

To specify pre-freeze and post-thaw scripts for the job:

1. Switch to the **Scripts** tab.
2. In the **Script processing mode** section, choose a scenario for script execution:
   - Select the **Require successful script execution** option if you want Veeam Backup & Replication to stop the backup process if the script fails.
   - Select the **Ignore script execution failures** option if you want to continue the backup process, even if script errors occur.
   - Select the **Disable script execution** option if you do not want to run scripts for the VM.
3. In the **Windows scripts** section, specify paths to pre-freeze and post-thaw scripts for Microsoft Windows VMs. For the list of supported script formats, see the Veeam Backup & Replication User Guide, section Pre-Freeze and Post-Thaw Scripts.
4. In the **Linux scripts** section, specify paths to pre-freeze and post-thaw scripts for Linux VMs. For the list of supported script formats, see the Veeam Backup & Replication User Guide, section Pre-Freeze and Post-Thaw Scripts.

> **TIP**
>
> If you have added a resource pool, hot or cluster with Microsoft Windows and Linux VMs to the job, you can select to execute both Microsoft Windows and Linux scripts for the VM container. When the job starts, Veeam Backup & Replication will automatically determine what OS type is installed on the VM and use the required scripts to quiesce this VM.

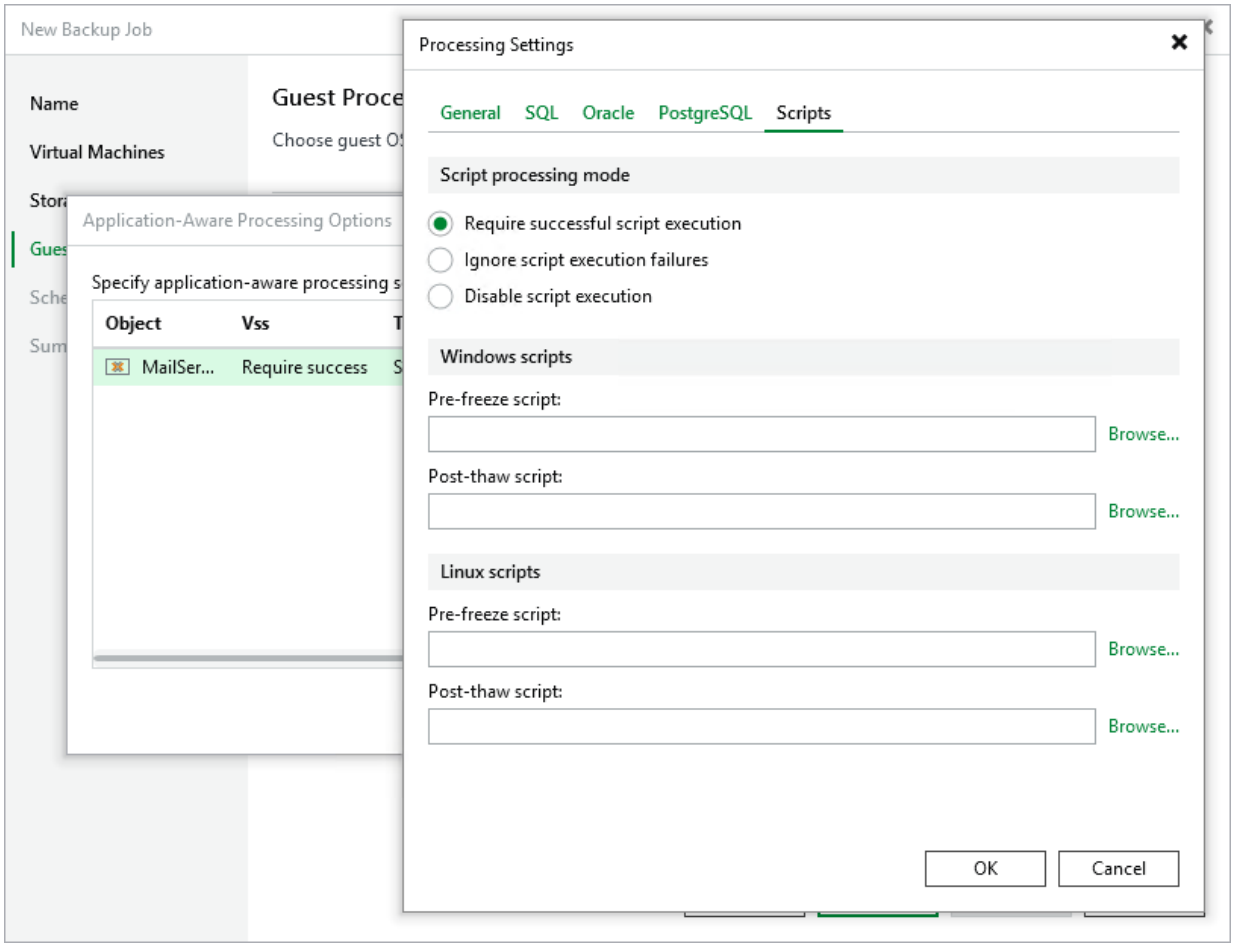
New Backup Job
Name
Virtual Machines
Guest Proce
Choose guest O
Application-Aware Processing Options
Specify application-aware processing s
Object
Vss
MailSer...
Require success
Processing Settings
General
SQL
Oracle
PostgreSQL
Scripts
Script processing mode
Require successful script execution
Ignore script execution failures
Disable script execution
Windows scripts
Pre-freeze script:
Browse...
Post-thaw script:
Browse...
Linux scripts
Pre-freeze script:
Browse...
Post-thaw script:
Browse...
OK
Cancel

Step 5b. Enable VM Guest OS File Indexing

To be able to recover individual files with 1 click and to search for specific items in Veeam Backup Enterprise Manager during file-level restore, you must enable file indexing to instruct Veeam Backup & Replication to create a catalog of files and folders that belong to VMs included into the backup scope. To do that, select the **Enable guest file system indexing and malware detection** check box at the **Guest Processing** step of the wizard.

> **NOTE**
>
> If you enable file indexing, Veeam Backup & Replication will scan VM data for suspicious file system activity and malware file presence every time the backup job completes successfully. For more information, see the Veeam Backup & Replication User Guide, section How Guest Indexing Data Scan Works.

By default, Veeam Backup & Replication will create a catalog of all files and folders for each processed VM — except for system files. To change this behavior and configure indexing settings for a specific VM, do the following:

1. Click **Customize guest file system indexing settings**.
2. In the **Guest File System Indexing Options** window, select the necessary VM and click **Edit** > **Windows indexing** or **Linux indexing**. You can configure indexing settings for one or more VMs at a time.
3. In the **Indexing Settings** window, choose whether you want to index files in all guest OS folders, to index files only in specific folders, or not to index any files at all.

   If you select the **Index everything except** or **Index only following folders** option, you will be able to modify the list of folders included into the indexing scope — either manually or by using system environment variables (for example, *%windir%*, *%ProgramFiles%* and *%Temp%*).

> **IMPORTANT**
>
> To allow Veeam Backup & Replication to perform guest OS file indexing for Linux VMs, `openssh`, `gzip` and `tar` utilities must be installed on the processed VMs.

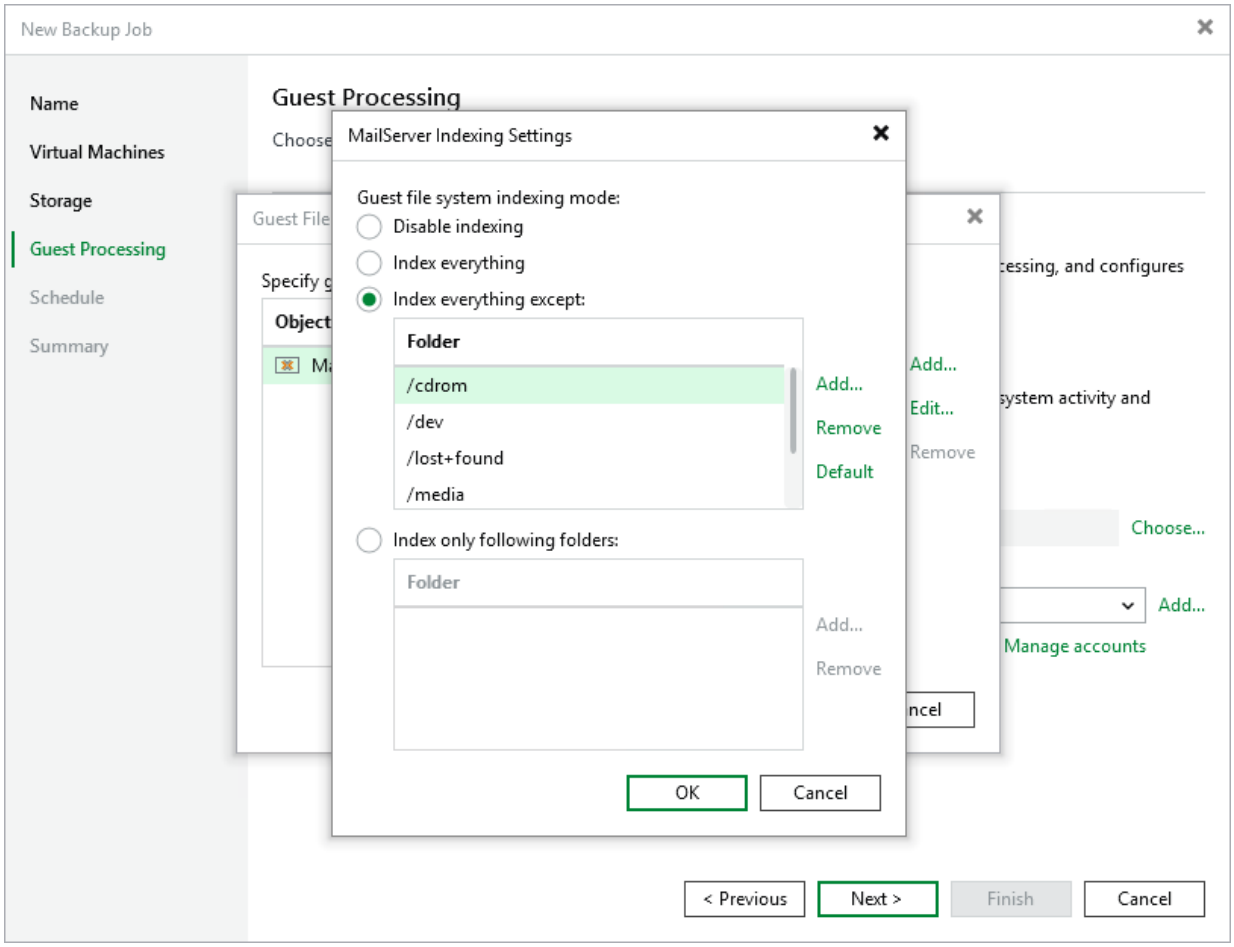
New Backup Job
Name
Virtual Machines
Storage
Guest Processing
Schedule
Summary
Guest Processing
MailServer Indexing Settings
Guest file system indexing mode:
Disable indexing
Index everything
Index everything except:
Folder
/cdrom
/dev
/lost+found
/media
Add...
Remove
Default
Index only following folders:
Folder
Add...
Remove
OK
Cancel
Add...
Edit...
Remove
Choose...
Add...
Manage accounts
< Previous
Next >
Finish
Cancel

Step 5c. Choose Guest Interaction Proxy

To produce transactionally consistent backups and to perform file system indexing, Veeam Backup & Replication communicates with the guest OS of each processed VM to deploy non-persistent runtime components that coordinate guest processing activities such as accessing VM applications and creating a catalog of VM files. Since these activities may significantly increase the load on the backup server in case of a large backup scope, Veeam Backup & Replication distributes the load among all Microsoft Windows and Linux servers added to the backup infrastructure (further referred to as guest interaction proxies).

By default, Veeam Backup & Replication automatically chooses which guest interaction proxy to use for each of the processed VMs based on network settings and rules listed in the Veeam Backup & Replication User Guide, section Guest Interaction Proxies. You can also manually limit the list of servers that may be used as proxies — to do that, click **Choose**, select the **Prefer the following guest interaction proxy servers** option and then select check boxes next to the necessary servers.

For a server to be displayed in the list of available log shipping servers, it must be added to the backup infrastructure as described in the Veeam Backup & Replication User Guide, sections Adding Microsoft Windows Servers and Adding Linux Servers.

> **IMPORTANT**
>
> Due to technical limitations, Linux-based proxies cannot access Windows guest OSes in the current version. That is why if you have added Windows-based VMs to the backup scope at step 3 of the wizard, you must also add at least one Microsoft Windows server to the backup infrastructure.

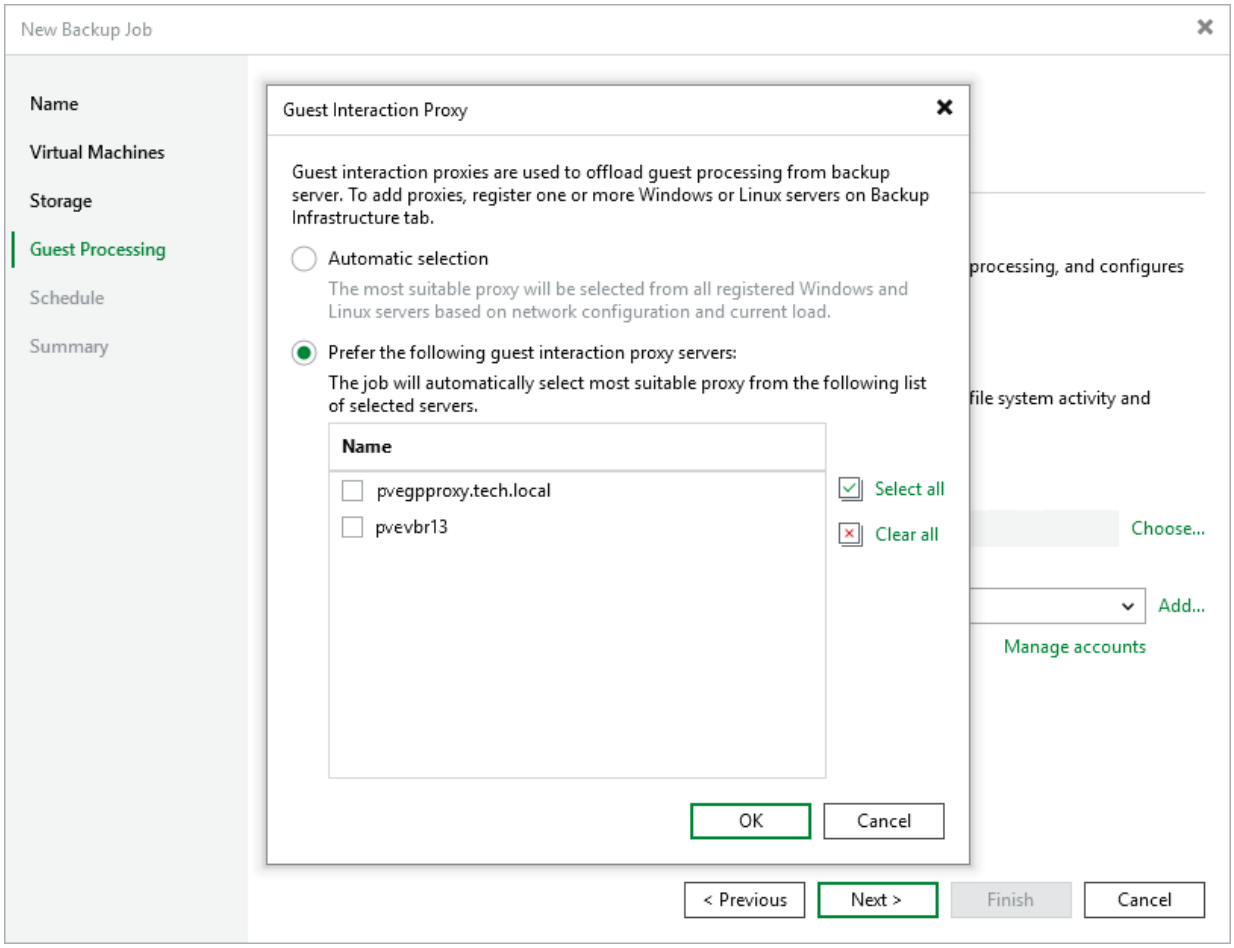

Step 5d. Manage VM Guest OS Credentials

If you enable application-aware processing or instruct Veeam Backup & Replication to create a catalog of VM files and folders, you must also specify a user whose credentials will be used to communicate with VM guest OSes. Note the specified user must have the permissions required to perform guest processing. For more information on the required permissions, see Planning and Preparation.

By default, Veeam Backup & Replication uses single credentials to access guest OSes of all VMs included into the backup scope. However, since Windows-based VMs and Linux-based VMs require different types of access credentials, you may need to specify the credentials explicitly for each of the processed VMs. To do that, click **Credentials**, select a VM in the **Guest OS credentials** window, and then click **Set User** > **Standard credentials** (for a Windows-based VM) or **Set User** > **SSH credentials** (for a Linux-based VM).

For a user to be displayed in the **Credentials** list, it must be added to the Credentials Manager as described in the Veeam Backup & Replication User Guide, section Credentials Manager. If you have not added the necessary user to the Credentials Manager beforehand, you can do it without closing the **New Job** wizard. To do that, click either the **Manage accounts** link or the **Add** button, and specify the user name, password and description in the **Credentials** window.

> **TIP**
>
> If the backup scope includes a resource pool, host or cluster, you can specify both Standard and SSH credentials. This will allow Veeam Backup & Replication to access the processed VMs regardless of their guest OSes.

To check whether Veeam Backup & Replication is able to connect to the VM guest OSes using the specified credentials, click **Verify Network Connectivity**.

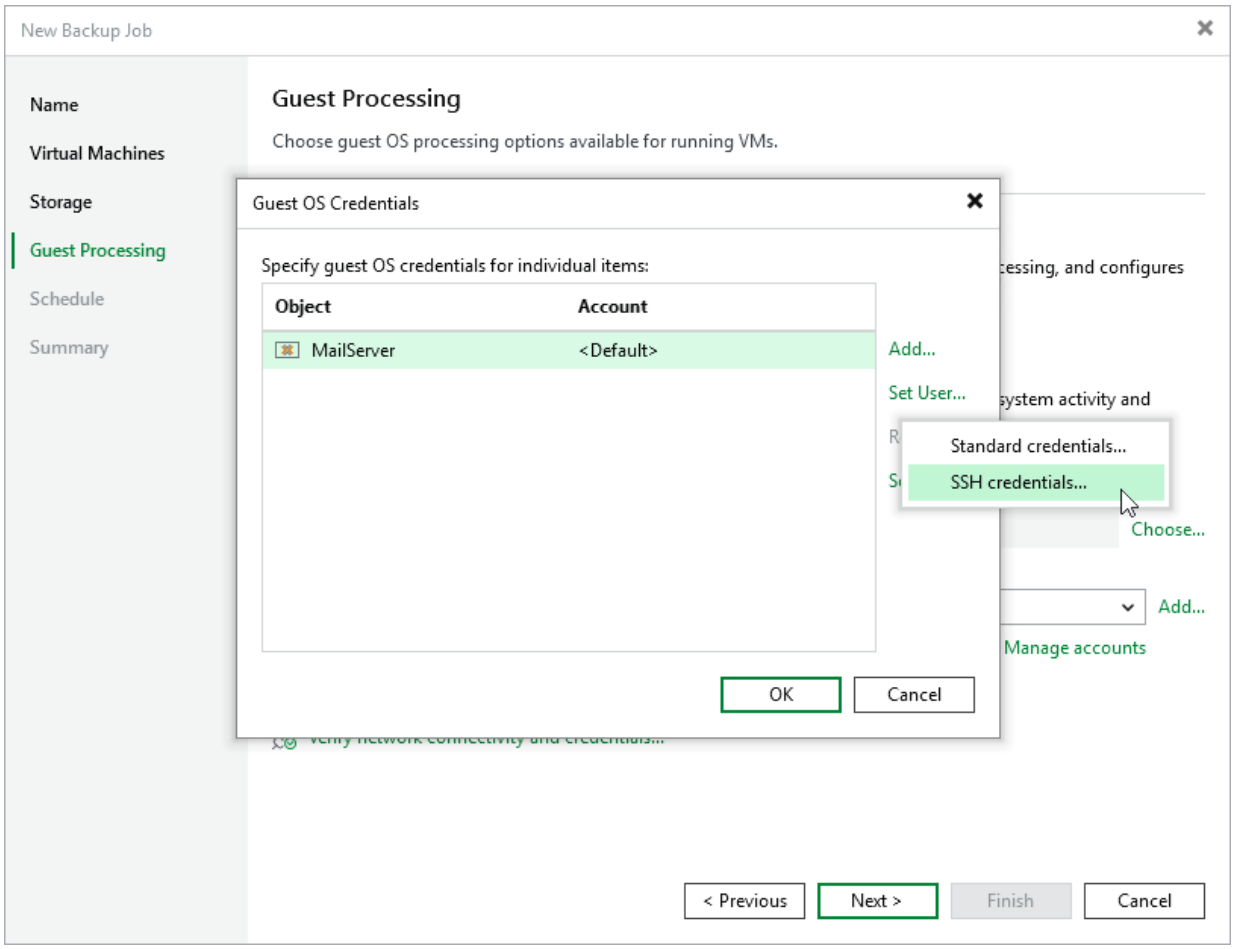

Step 6. Specify Job Scheduling Options

At the **Schedule** step of the wizard, you can instruct Veeam Backup & Replication to start the backup job automatically according to a specific backup schedule. The backup schedule defines how often data of the VMs added to the backup job will be backed up.

Veeam Backup & Replication allows you to create schedules of the following types:

- **Daily at this time** – the backup job will create restore points at a specific time on specific days.
- **Monthly at this time** – the backup job will create restore points once a month on a specific day.
- **Periodically every** – the backup job will create restore points repeatedly with a specific time interval every day.

> **TIP**
>
> You can instruct Veeam Backup & Replication to run the backup job again if it fails on the first try. To do that, select the **Retry failed items processing** check box, and specify the maximum number of attempts to run the backup job and the time interval between retries. When retrying backup jobs, Veeam Backup & Replication processes only those VMs that failed to be backed up during the previous attempt.

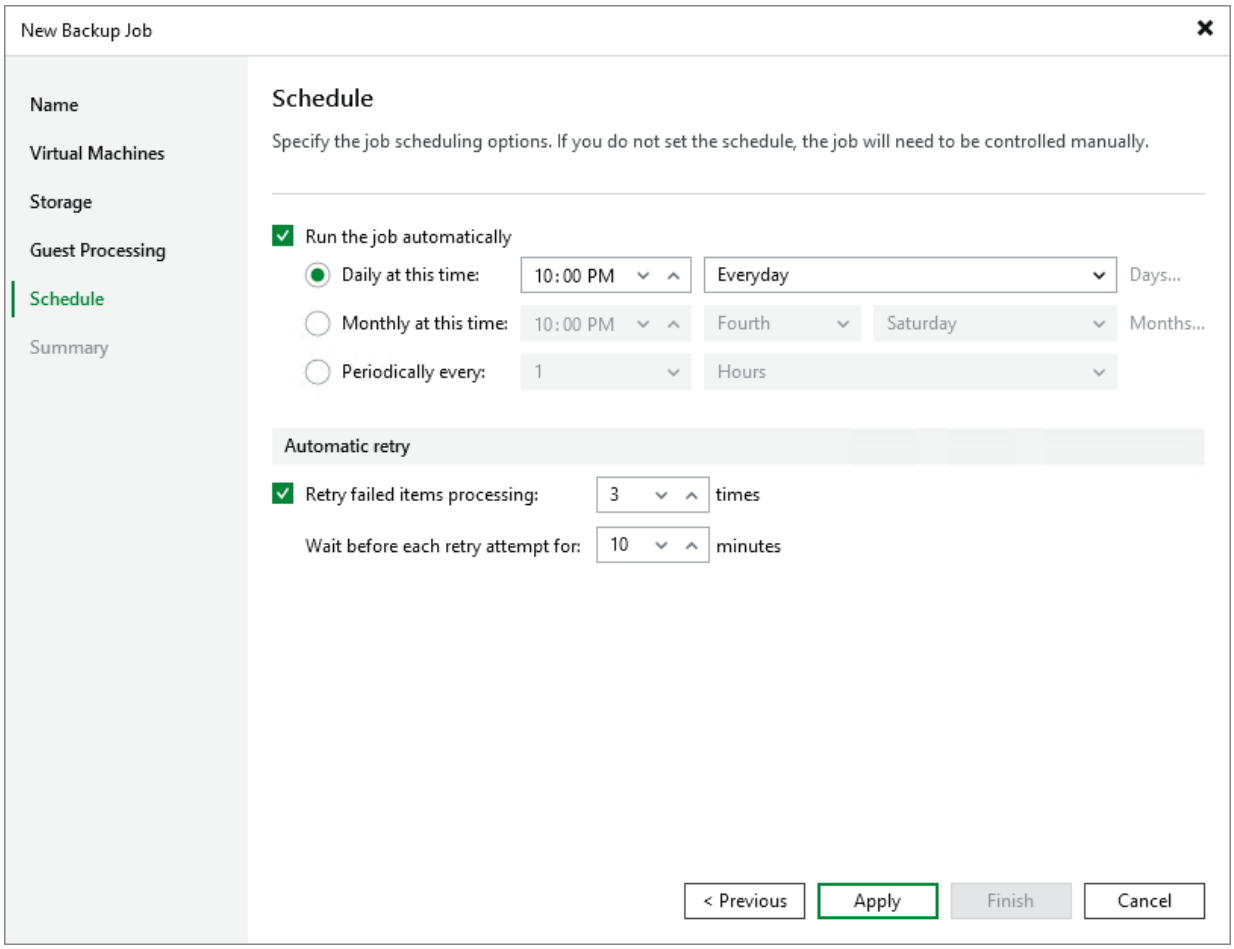

Step 7. Finish Working with Wizard

At the **Summary** step of the wizard, review summary information and click **Finish**. As soon as Veeam Backup & Replication starts the job, the backup progress will be displayed in the working area when you navigate to **Jobs** > **Backups** in the inventory pane of the **Home** view.

> **TIP**
>
> If you want to start the job immediately, select the **Run the job when I click Finish** check box and then click **Finish**.

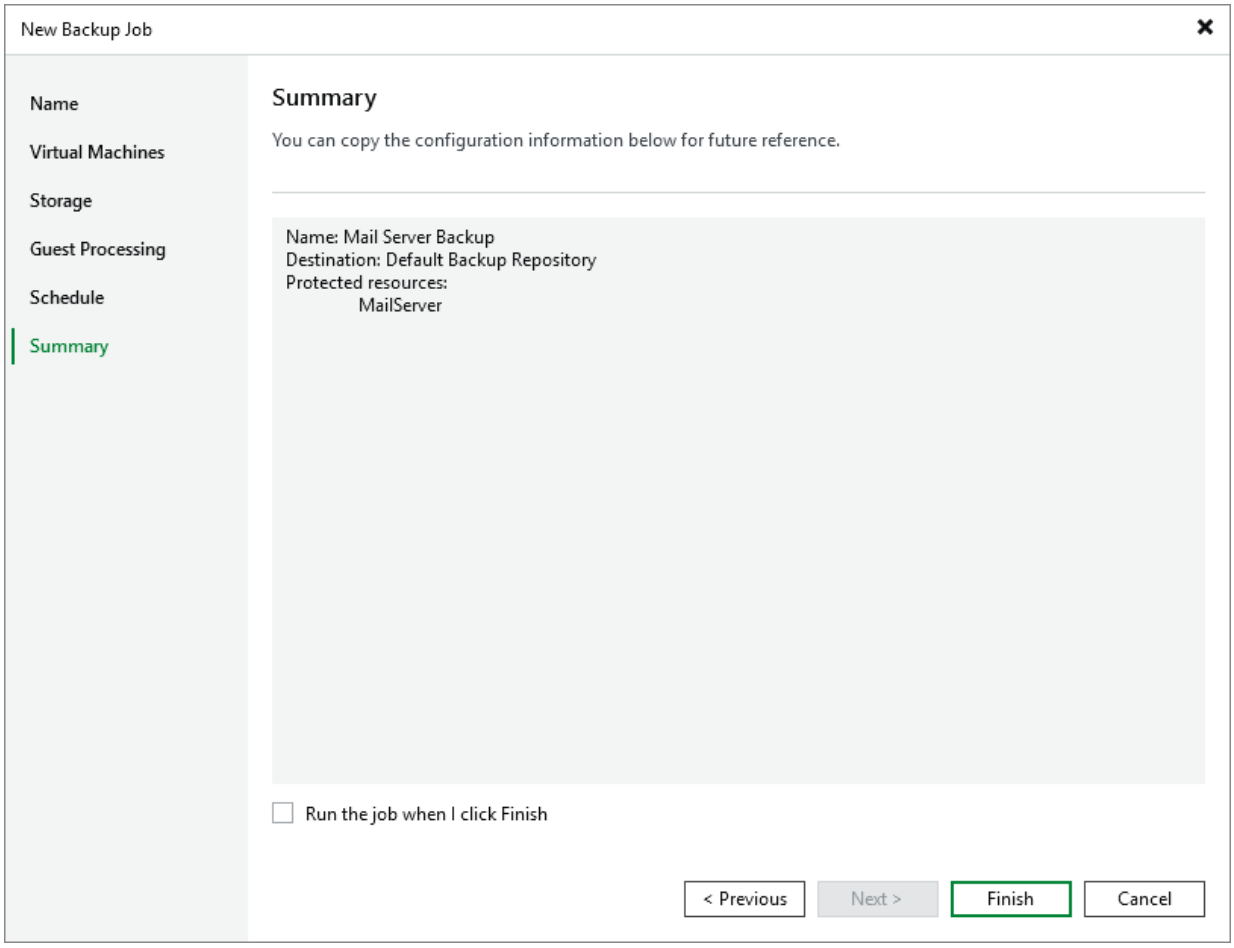

Starting and Stopping Backup Jobs

You can start a backup job manually, for example, if you want to create an additional restore point and do not want to modify the configured job schedule. You can also stop a backup job manually if processing of an Proxmox VE VM is about to take too long, and you do not want the job to have an impact on the production environment during business hours. When you stop a running job, Veeam Backup & Replication creates a new restore point only for those VMs that have already been processed by the time you stop the job.

To start or stop a backup job, do the following:

1. Open the **Home** view.
2. In the inventory pane, select **Jobs**.
3. In the working area, right-click the necessary job and select **Start** or **Stop**.

   Alternatively, select the job and click **Start** or **Stop** on the ribbon.

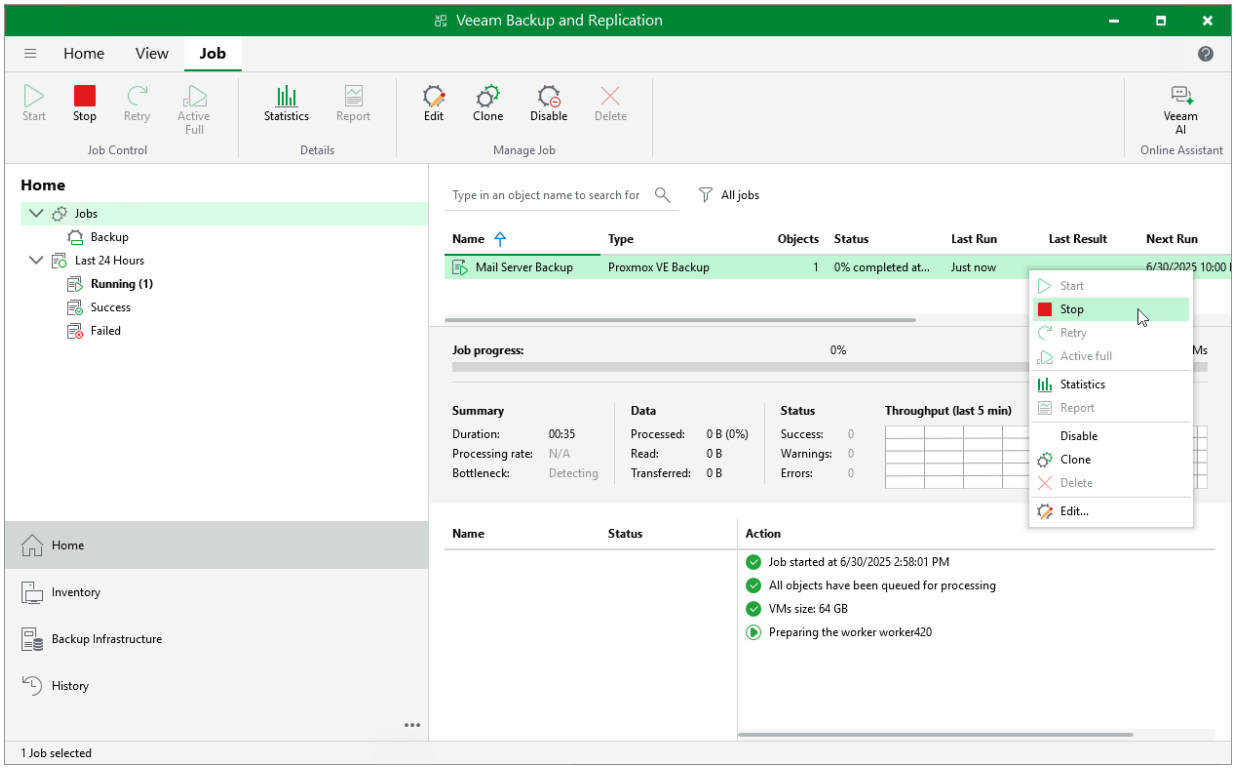

Retrying Jobs

If a job fails, you can retry the backup operation. When you perform a retry, Veeam Backup & Replication restarts the operation only for the failed resources added to the job and does not process VMs that have been processed successfully. As a result, retrying a job takes less time compared to restarting the job for all resources.

To retry a job, do the following:

1. Open the **Home** view.
2. In the inventory pane, select **Jobs**.
3. In the working area, right-click the necessary job and select **Retry**.

   Alternatively, select the job and click **Retry** on the ribbon.

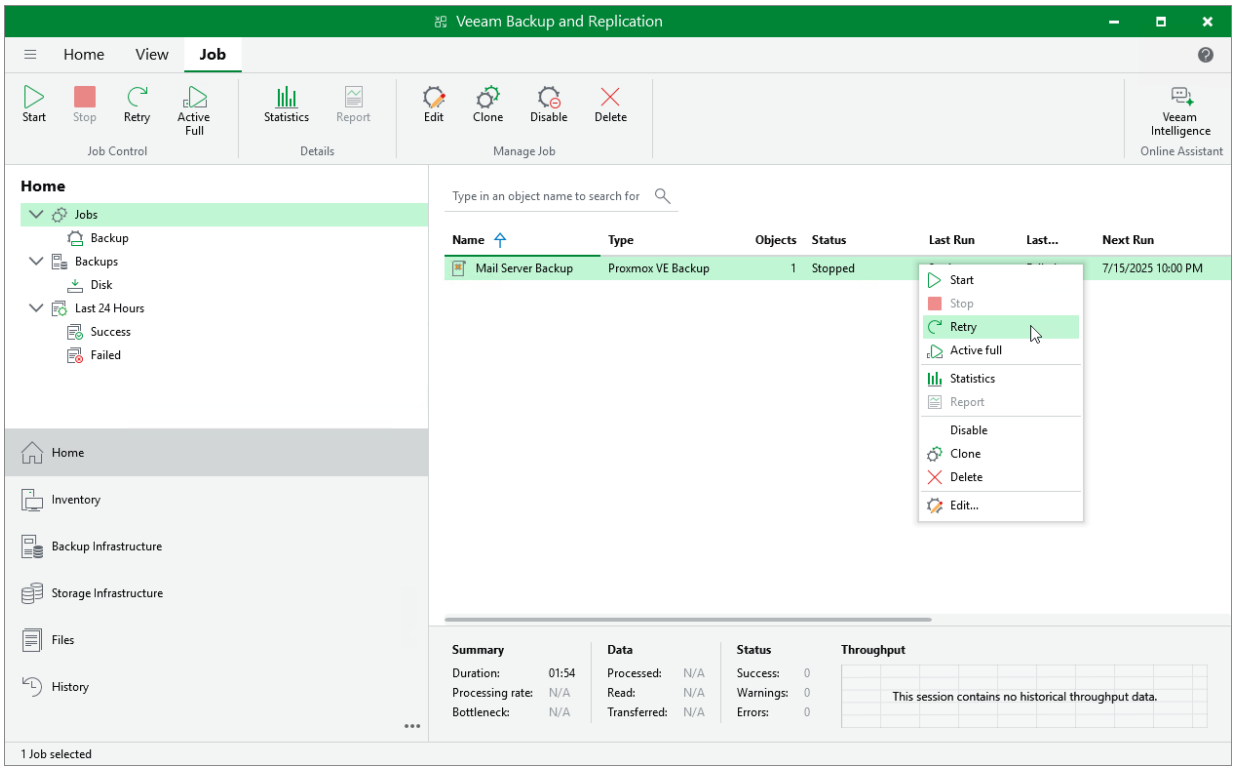

Editing Backup Job Settings

For each backup job, you can modify settings configured while creating the job.

1. Open the **Home** view.
2. In the inventory pane, select **Jobs**.
3. In the working area, right-click the job and select **Edit**.

   Alternatively, select the necessary job and click **Edit** on the ribbon.
4. Complete the **Edit Job** wizard:
   a. To provide a new name and description for the job, follow the instructions provided in section Creating Backup Jobs (step 2).
   b. To edit the backup scope, follow the instructions provided in section Creating Backup Jobs (step 3).
   c. To change the backup repository where backups are stored, to configure backup job retention settings, to schedule active and synthetic full backups, to configure health checks and email notifications, follow the instructions provided in section Creating Backup Jobs (step 4).
   d. To modify settings for application-aware processing of VMs included into the backup scope, follow the instructions provided in section Creating Backup Jobs (step 5).
   e. To modify the job schedule and configure automatic retry settings, follow the instructions provided in section Creating Backup Jobs (step 6).
   f. At the **Summary** step of the wizard, review configuration information and click **Finish**.

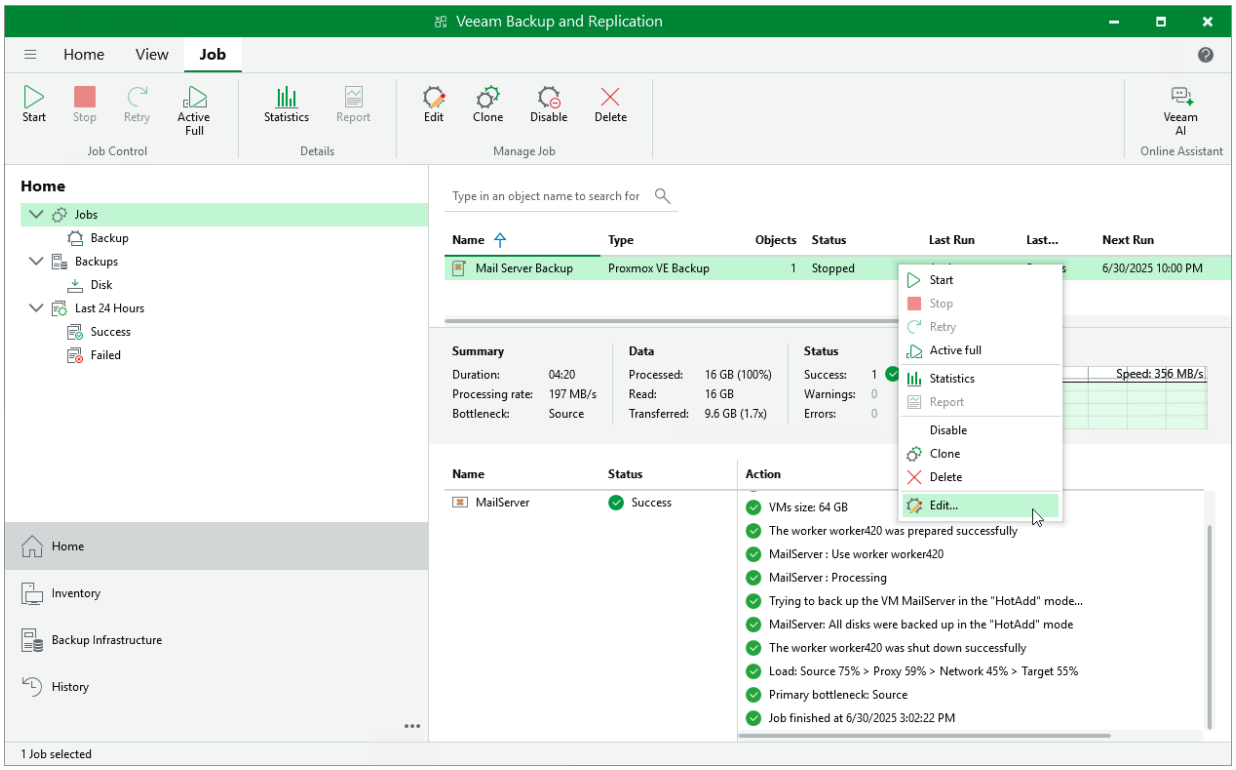

Analyzing Performance Bottlenecks

As any backup application handles a great amount of data, it is important to make sure the data flow is efficient and all resources engaged in the backup process are optimally used. For backup jobs, Veeam provides advanced statistics about the data flow efficiency and lets you identify bottlenecks at the following stages of the data transmission process:

1. Reading VM data blocks from the source.
2. Processing VM data on a worker.
3. Transporting data over the network.
4. Writing data to the target.

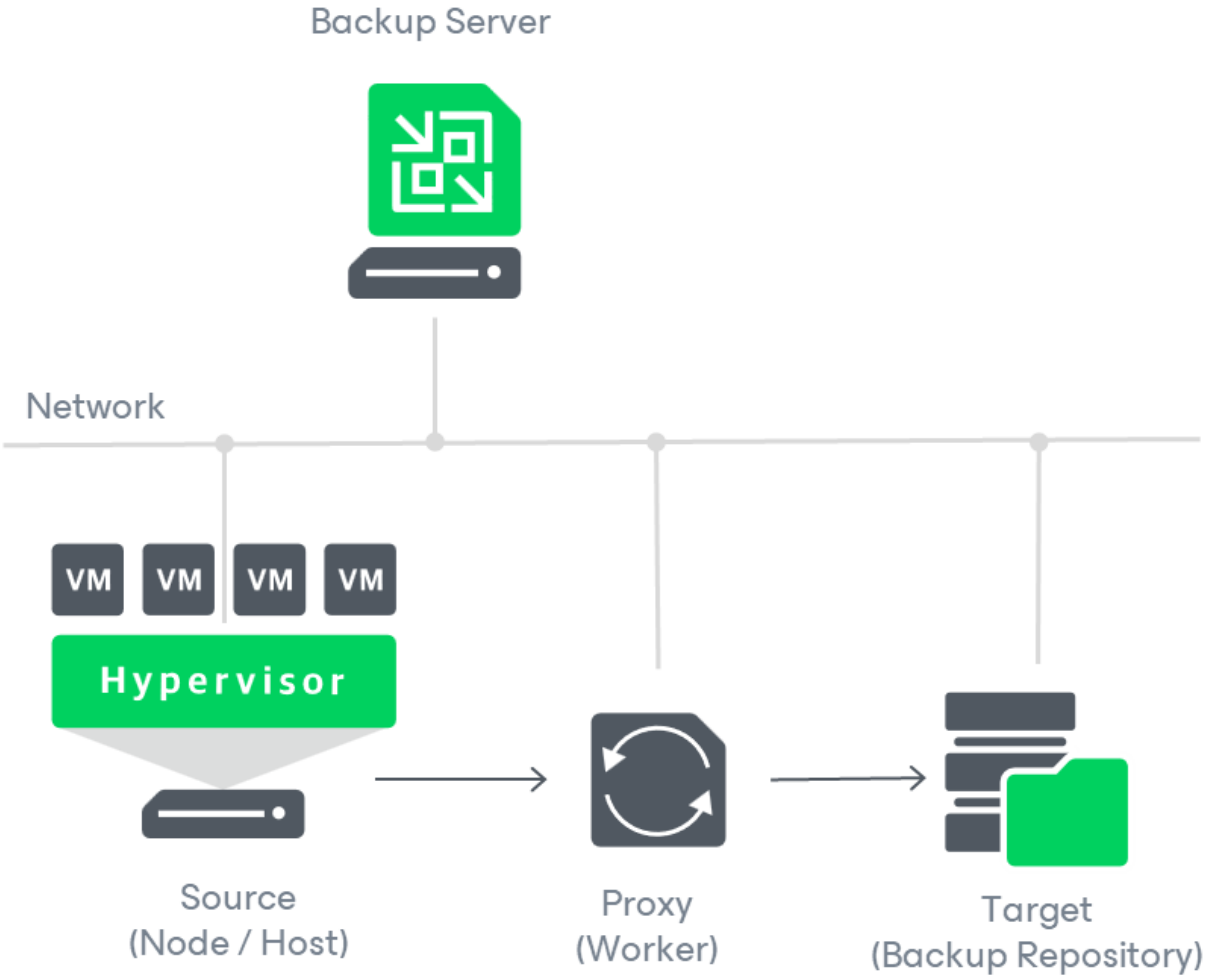

While evaluating the data transmission process, Veeam Backup & Replication analyzes performance of all the data flow components:

- **Source** – the source disk reader component responsible for retrieving data from the source node.
- **Proxy** – the worker component responsible for processing VM data.
- **Network** – the network queue writer component responsible for getting processed VM data from the worker and sending it over the network to the Target (directly or through the Gateway Server).

- **Target** – the gateway server component responsible for processing VM data, or the target disk writer component responsible for storing data in the backup repository.

To see the bottleneck statistics for a job or a specific VM processed by the job, do the following:

1. Open the **Home** view.
2. In the inventory pane, select **Jobs**.
3. In the working area, right-click the backup job for which you want to see the bottleneck statistics and select **Statistics**.

   Alternatively, select the job and click **Statistics** on the ribbon.
4. In the job session details window, check the **Bottleneck** field in the **Summary** column.

> **TIP**
>
> To see the bottleneck statistics for a specific VM, select the VM name in the **Name** column and check the **Load** record in the **Action** column. To learn how to analyze the statistics, see Veeam Backup & Replication User Guide, section Performance Bottlenecks.

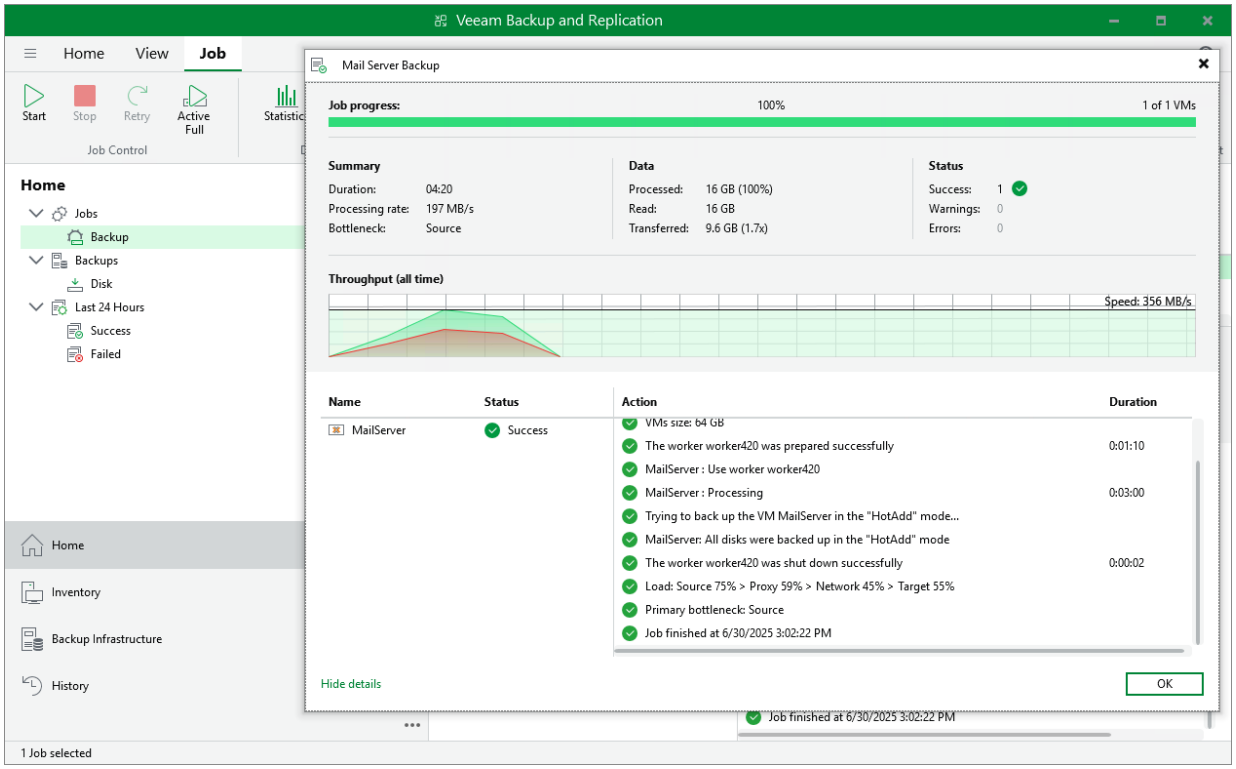

Cloning Backup Jobs

You can create a new job by cloning an existing one. Job cloning allows you to create an exact copy of any job with the same job settings.

To clone a job, do the following:

1. Open the **Home** view.
2. In the inventory pane, select **Jobs**.
3. In the working area, right-click the necessary job and select **Clone**.

   Alternatively, select the necessary job and click **Clone** on the ribbon.

The name of the cloned job is formed by the following rule: *<job_name_clone1>*, where *job_name* is the name of the original job and *clone1* is a suffix added to the original job name. If you clone the same job again, the number in the name will be incremented, for example, *job_name_clone2*, *job_name_clone3* and so on. To change the name of a cloned job, edit the job as described in section Editing Backup Job Settings.

> **NOTE**
>
> If the original job is scheduled to run automatically, Veeam Backup & Replication disables the cloned job. To enable the cloned job, select it in the job list and click **Enable**.

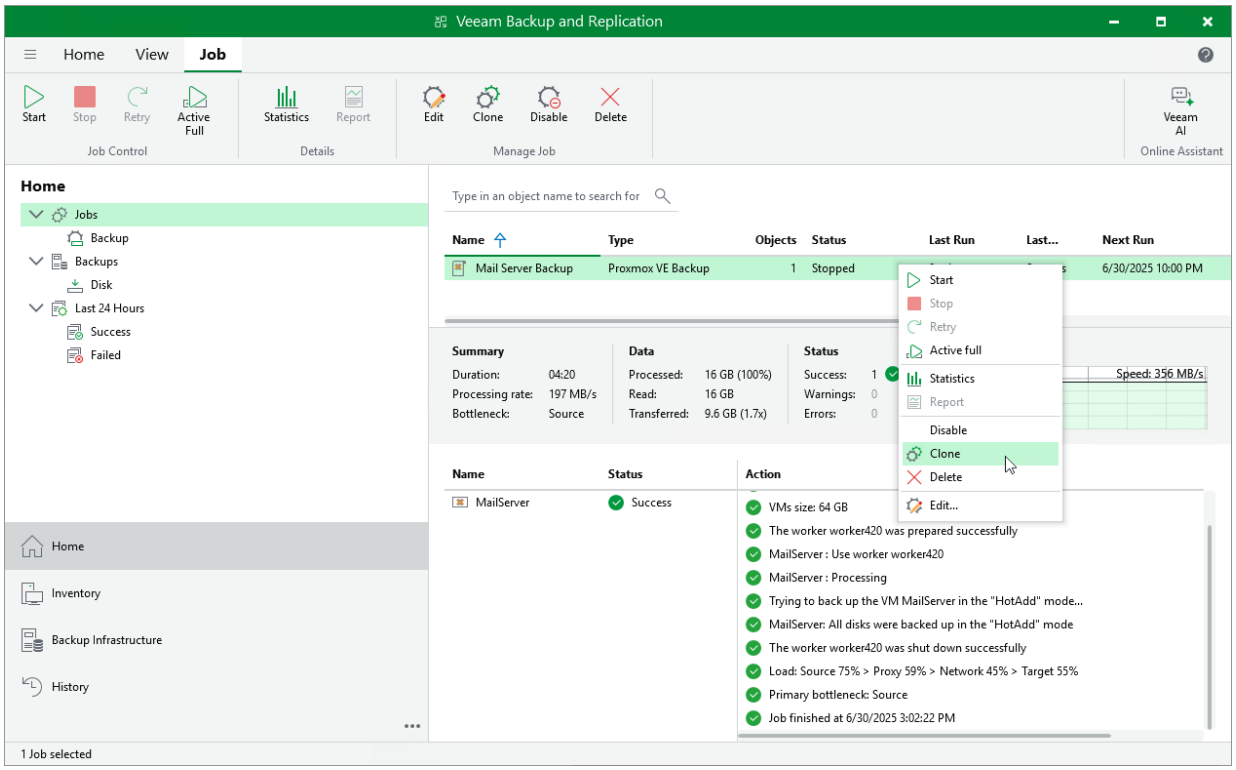

Enabling and Disabling Backup Jobs

By default, all created backup jobs run according to the specified schedules. However, you can temporarily disable a job so that it does not run automatically. You will still be able to enable the disabled job at any time you need.

To enable or disable a backup job, do the following:

1. Open the **Home** view.
2. In the inventory pane, select **Jobs**.
3. In the working area, right-click the necessary job and select **Disable**.

   Alternatively, select the job and click **Disable** on the ribbon.

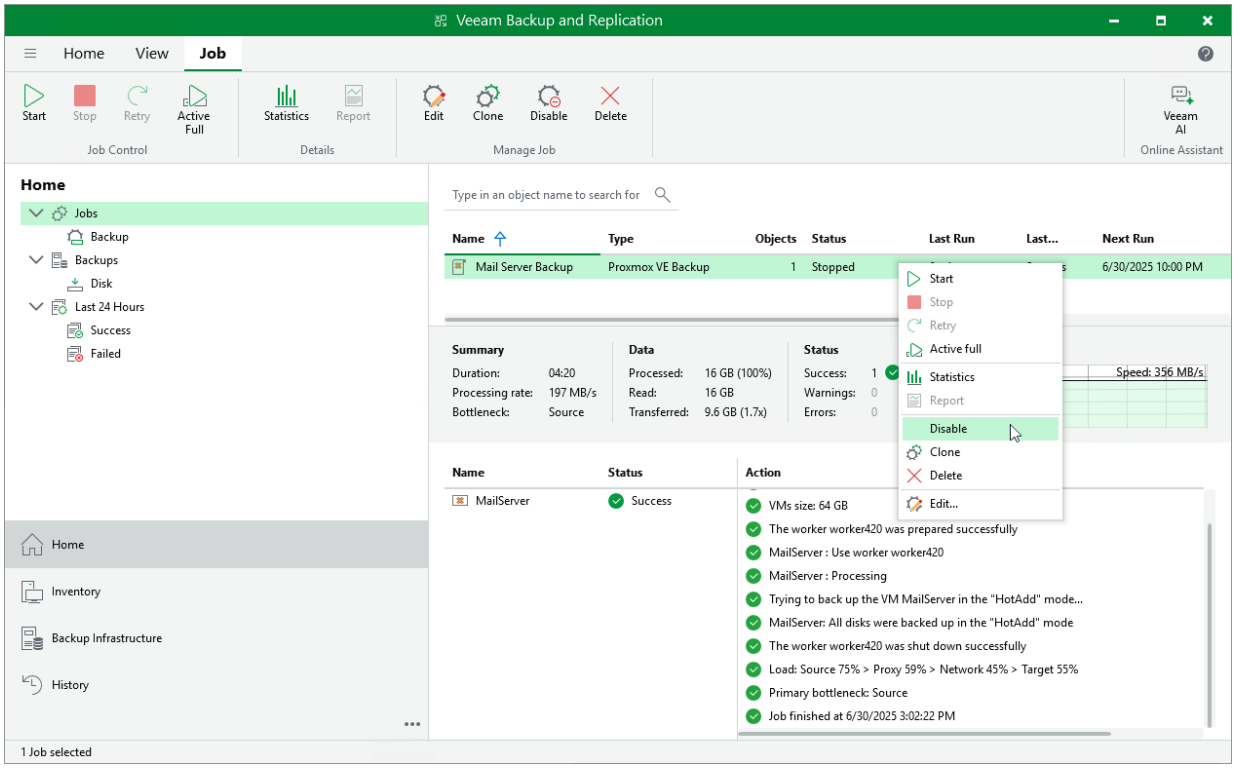

Deleting Backup Jobs

You can permanently delete a backup job from the Veeam Backup & Replication configuration database if you no longer need it. When you delete a job, backups created by this job are displayed under the **Backups** > **Disk (Orphaned)** node in the **Home** view of the Veeam Backup & Replication console. If you want to delete backup files as well, follow the instructions provided in section Deleting Backups.

To delete a backup job, do the following:

1. Open the **Home** view.
2. In the inventory pane, select **Jobs**.
3. In the working area, right-click the necessary job and select **Delete**.

   Alternatively, select the job and click **Delete** on the ribbon.

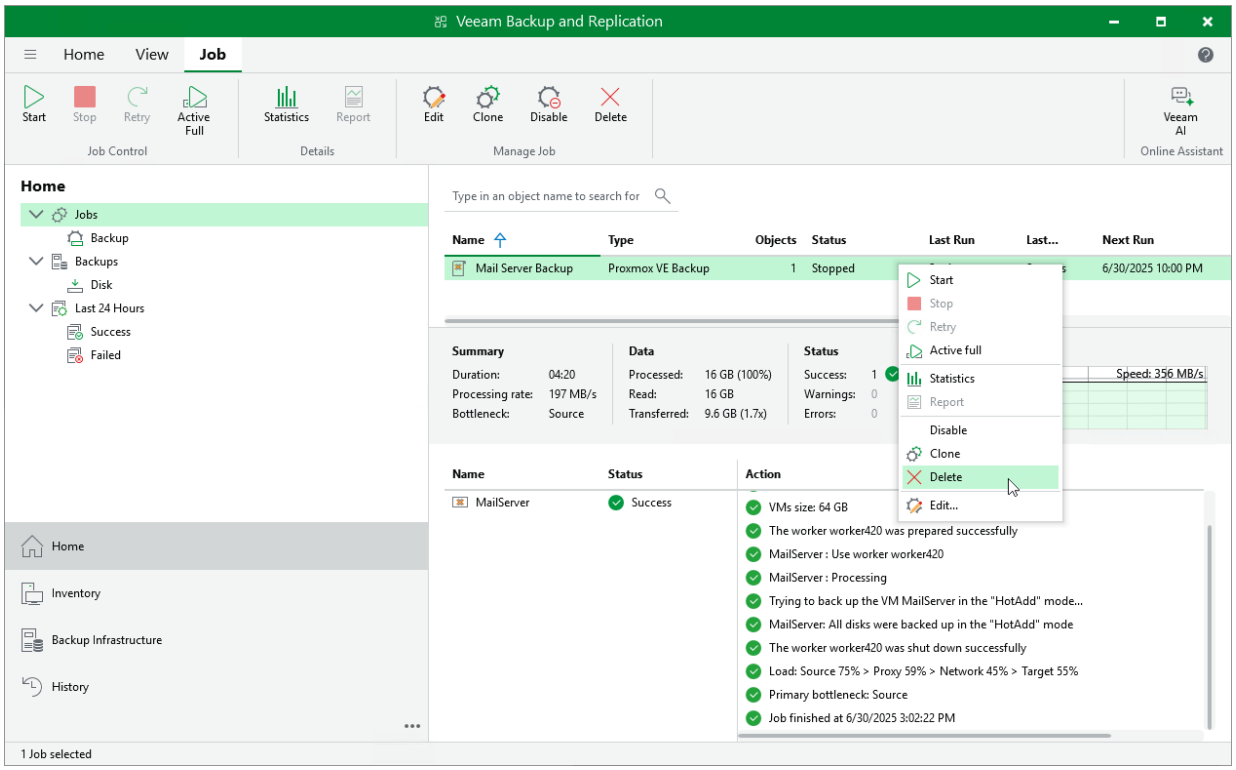

Creating Active Full Backups

You can manually create an active full backup for all VMs added to a backup job:

1. Open the **Home** view.
2. In the inventory pane, select **Jobs**.
3. In the working area, right-click the necessary job and select **Active full**.

   Alternatively, select the job and click **Active Full** on the ribbon.

> **TIP**
>
> To create active full backup automatically according to a specific schedule, configure backup job settings as described in section Creating Backup Jobs (step 4).

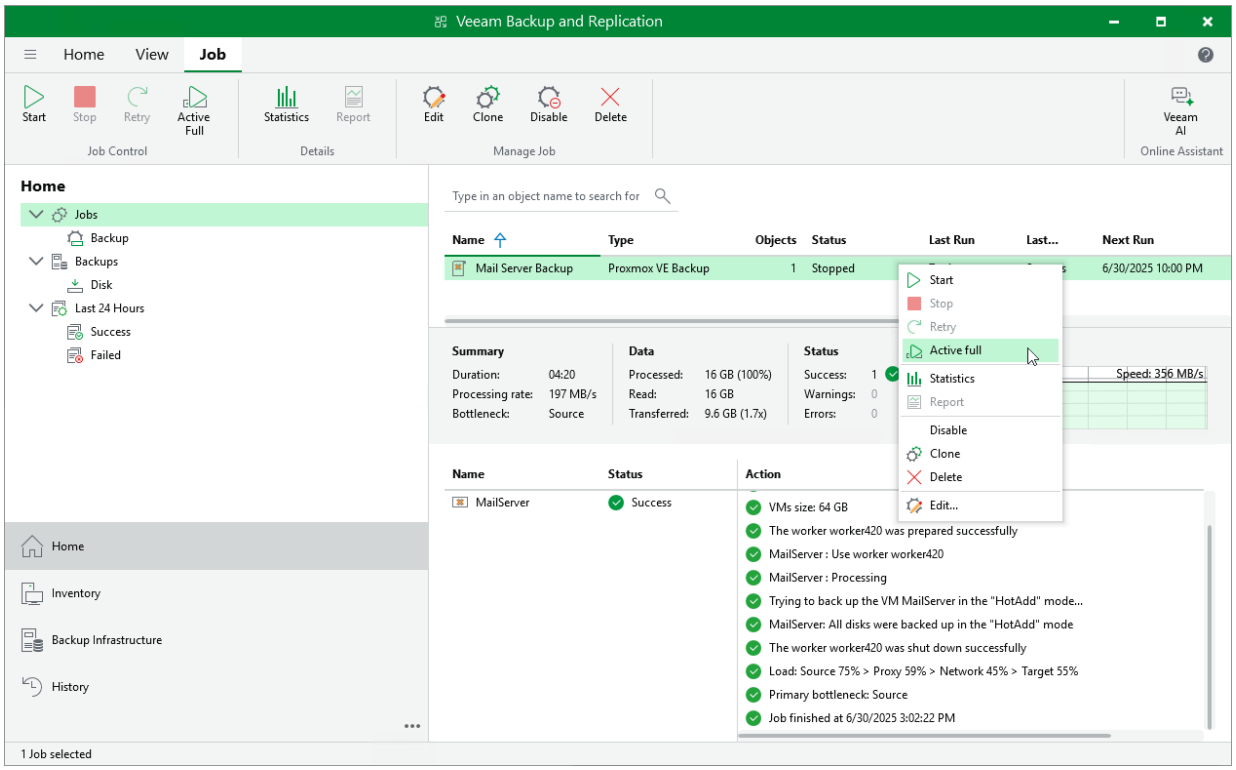

Creating VeeamZIP Backups

You can back up one or multiple VMs without configuring backup jobs. To do that, you can leverage the VeeamZIP feature — it can be helpful, for example, if you want to create backups for VMs immediately, archive VMs before decommissioning and so on. VeeamZIP produces a full backup that acts as an independent restore point. You can store the backup in a repository added to the backup infrastructure, in a local folder on the backup server or in a network share.

NOTE

- You cannot store VeeamZIP backups in HPE Cloud Bank Storage repositories.
- Veeam Backup & Replication does not apply network traffic throttling rules to VeeamZIP backup sessions. For more information, see the Veeam Backup & Replication User Guide, section Configuring Network Traffic Rules.

To create a VeeamZIP backup, do the following:

1. Open the **Inventory** view.
2. In the inventory pane, select **Virtual Infrastructure** > **Proxmox VE**.
3. In the working area, right-click the VM that you want to back up and select **VeeamZIP**.

   Alternatively, select the VM and click **VeeamZIP** on the ribbon.
4. Select the destination where the VeeamZIP backup will be stored.

TIP

You cannot specify an SMB share that requires authentication as a local or shared folder. However, you can add the SMB share to the backup infrastructure and specify it as backup repository.

The created VeeamZIP backup will be displayed under the **Backups** > **Disk (Exported)** node in the **Home** view of the Veeam Backup & Replication console.

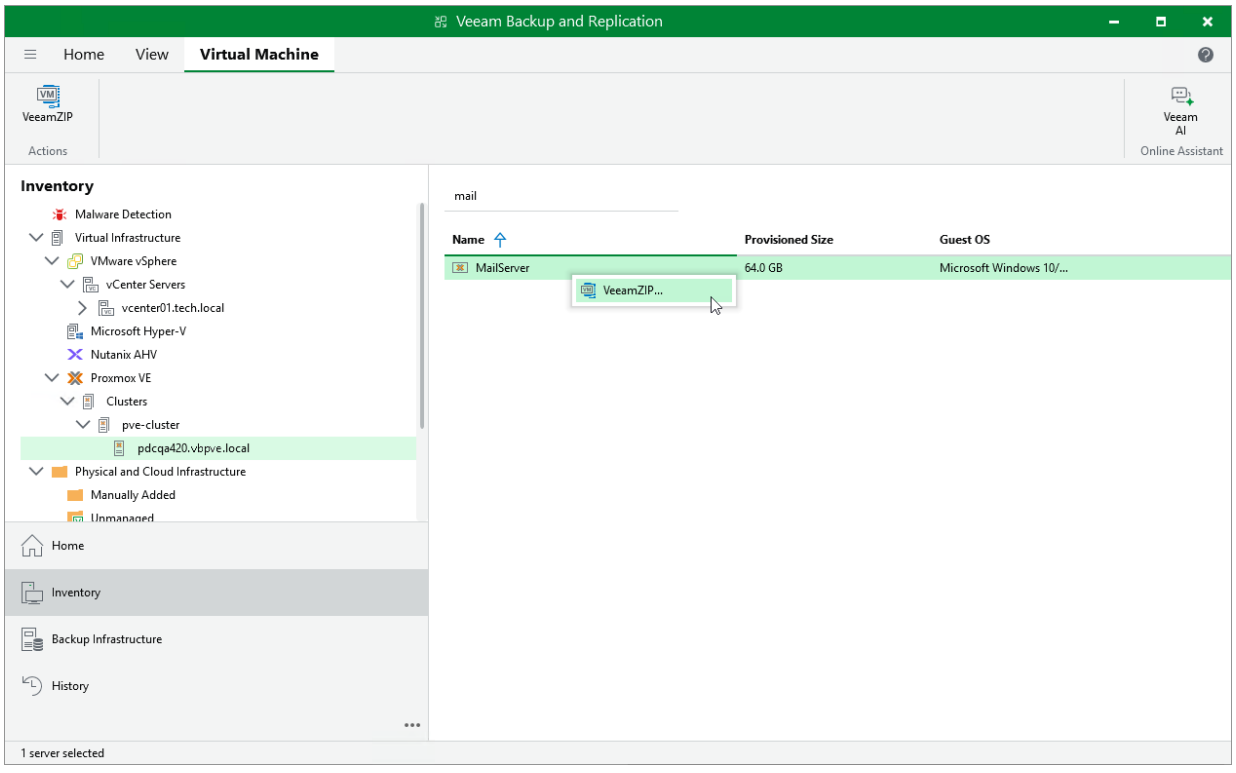

Managing Backups

Veeam Backup & Replication stores information on all protected Proxmox VE VMs in the configuration database. Even if a VM is no longer protected by any configured backup job and even if the VM no longer exists in the Proxmox VE environment, records about created backups will not be deleted from the database until Veeam Backup & Replication automatically removes all restore points associated with this VM according to the retention settings saved in the backup metadata. You can manage Proxmox VE VM backups as long as their records are present in the configuration database.

Viewing Backup Properties

After a backup job successfully creates a VM backup according to the specified schedule, or after you create an active full VM backup manually, the backup is displayed under the **Backups** node in the **Home** view of the Veeam Backup & Replication console. Each backup and the collection of restore points created for this backup is represented with a set of properties, such as:

- **Object Name** – the name of a protected VM.
- **Original Size** – the total amount of disk space allocated to the VM.
- **File Name** – the name of a restore point.
- **Data Size** – the amount of processed VM data.
- **Backup Size** – the amount of backed-up VM data.
- **Data Reduction** – the ratio between the data size and backup size.
- **Date** – the date and time when the restore point was created.
- **Immutable Until** – the date when the immutability period for the restore point will expire.
- **Status** – the result of the most recent malware scan performed for the restore point.

> **TIP**
>
> If a restore point is marked as *Infected* but you know that this point is clean, you can change its status manually. To do that, select the restore point and click **Malware** > **Mark as clean**. To learn how to manage infected restore points, see Veeam Backup & Replication User Guide, section Managing Malware Status.

To view backup properties, do the following:

1. Open the **Home** view.
2. In the inventory pane, select **Backups**.

3. In the working area, right-click the necessary backup job and select **Properties**.

   Alternatively, select the backup job and click **Properties** on the ribbon.

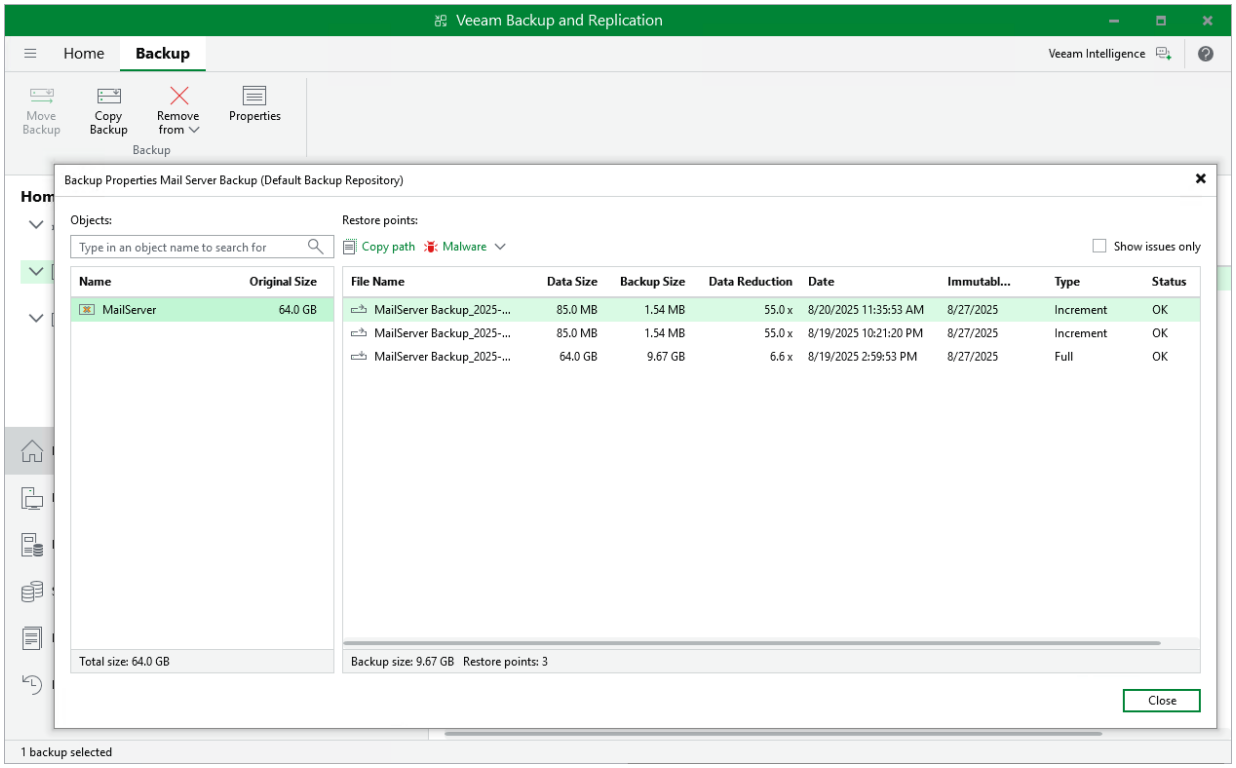

Verifying Backups

To perform an integrity check of VM backups, Veeam Backup & Replication offers the SureBackup technology that allows you to ensure that the created restore points are not corrupted. You can also scan the restore points with antivirus software installed on the backup server, and run YARA rules to detect malware and sensitive data.

To create a SureBackup job, do the following:

1. Open the **Home** view.
2. In the inventory pane, select **Jobs > Backup** and click **SureBackup Job** on the ribbon.
3. At the **Name** step of the **New SureBackup Job** wizard, select the **Backup verification and content scan only** verification mode, and then complete the wizard as described in the Veeam Backup & Replication User Guide, section Creating SureBackup Jobs.

If any of the verification checks fail for a restore point, Veeam Backup & Replication will mark both this restore point and all subsequent points in the backup chain as *Infected*. To learn how to manage infected restore points, see Veeam Backup & Replication User Guide, section Managing Malware Status.

> **TIP**
>
> You can scan backups of VMs manually on demand, without creating a SureBackup job. To learn how to do that, see the Veeam Backup & Replication User Guide, section Scan Backup.

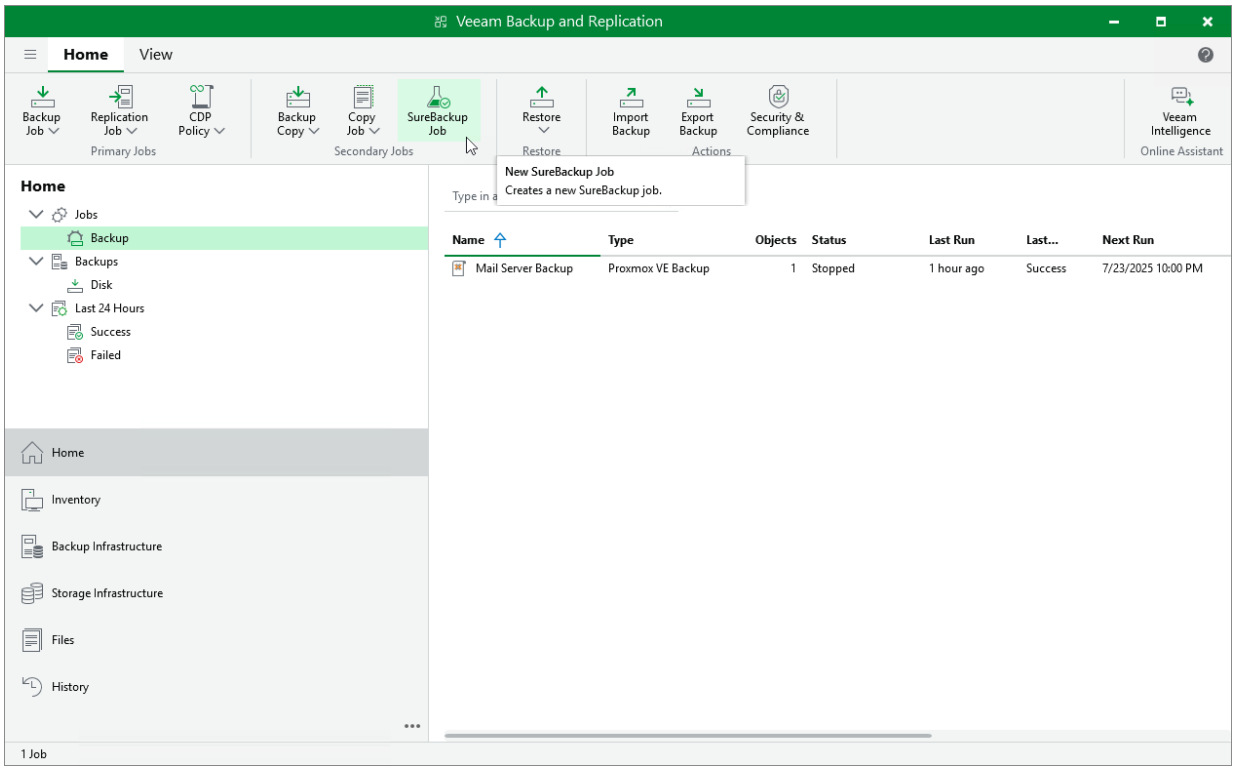

Exporting Backups

Exporting backups allows you to synthesize a complete and independent full backup file using restore points located in your backup repositories. That is, you can transform any backup chain into a standalone full backup file and save it to a repository or folder.

To export a backup, do the following:

1. Open the **Home** view.
2. In the inventory pane, select **Backups**.
3. In the working area, expand the job that created the backup, right-click the VM for which you want to synthesize a full backup file, and select **Export backup**.

   Alternatively, expand the job that created the backup, select the VM and click **Export Backup** on the ribbon.
4. Complete the **New Export** wizard as described in the Veeam Backup & Replication User Guide, section Performing Export.

Once the export operation completes, the exported backup will be displayed under the **Backups** > **Disk (Exported)** node in the **Home** view of the Veeam Backup & Replication console.

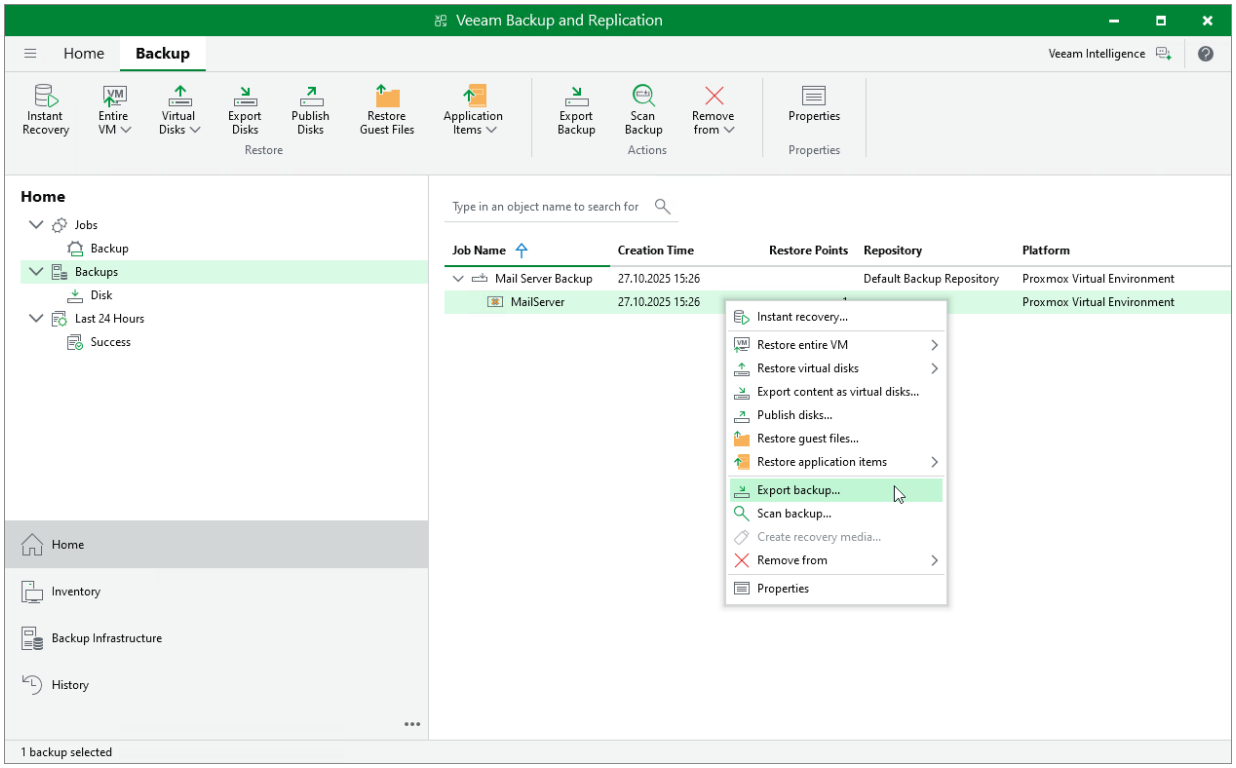

Copying Backups

With backup copy, you can create several instances of a backup and copy them to secondary (target) backup repositories for long-term storage. Target backup repositories can be located in the same site as the source backup repository or can be deployed off-site. Since the backup copy has the same format as the original backup, you can restore VM data directly from the backup copy in case a disaster strikes. For more information on the backup copy functionality, see the Veeam Backup & Replication User Guide, section Backup Copy.

To copy backups to a secondary backup repository, do the following:

1. Open the **Home** view.
2. In the inventory pane, select **Jobs > Backup** and click **Backup Copy > Image-level backup** on the ribbon.
3. Create a backup copy job as described in the Veeam Backup & Replication User Guide, section Creating Backup Copy Jobs.

Note that for backup copies, you can also use Veeam Cloud Connect repositories if a service provider is added to Veeam Backup & Replication.

> **TIP**
>
> Alternatively, you can create a copy of a backup without configuring a job as described in the Veeam Backup & Replication User Guide, section Copying Backups.

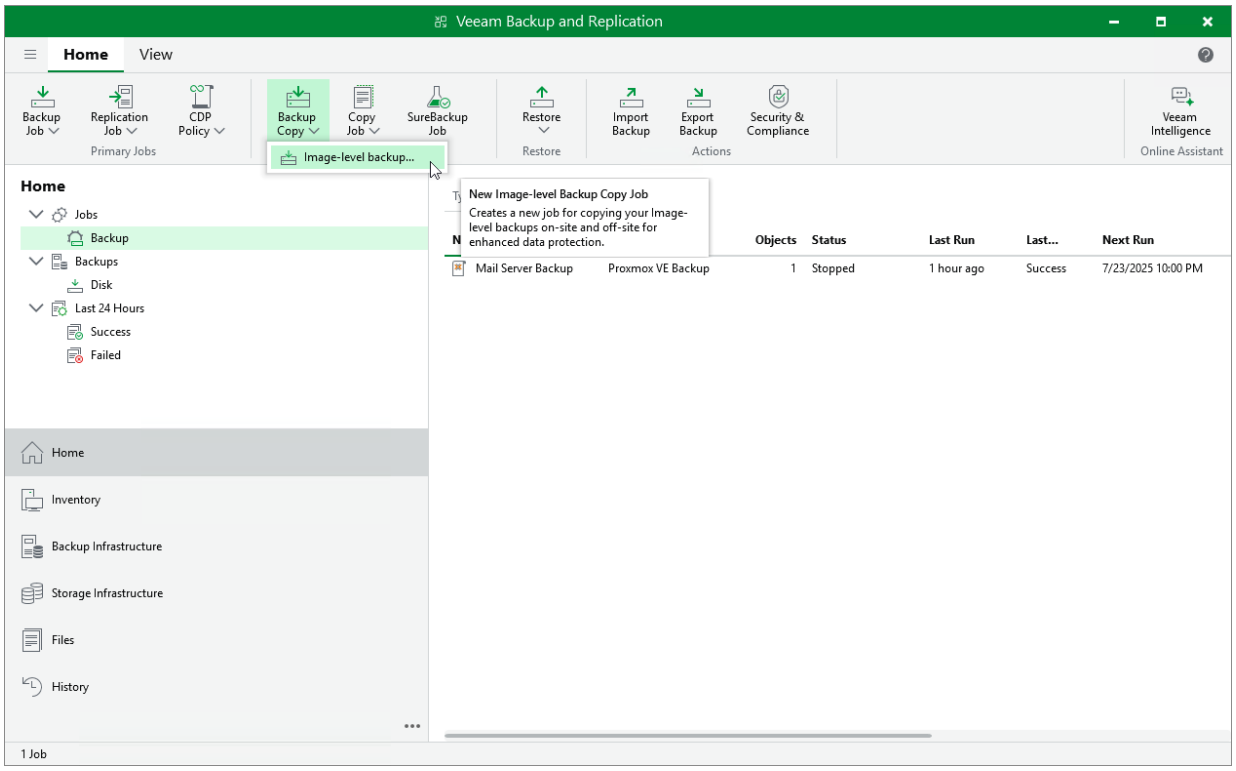

Copying Backups to Tapes

You can create archives of VM backups and copy them to tapes for long-term storage.
Veeam Backup & Replication allows you to manage tape archives the same way you manage backups in backup repositories. However, it usually takes more time to access archived data on tapes than to access backed-up data in repositories. For more information on tapes, see the Veeam Backup & Replication User Guide, section Tape Devices Support.

To archive VM backups to tape, do the following:

1. Configure the tape infrastructure:
    a. Connect tape devices as described in the Veeam Backup & Replication User Guide, section Tape Devices Deployment.
    b. Perform initial configuration of the tape infrastructure as described in the Veeam Backup & Replication User Guide, section Getting Started with Tapes (steps 1-3).
2. Create a backup to tape job as described in the Veeam Backup & Replication User Guide, section Creating Backup to Tape Jobs.

> **NOTE**
>
> You cannot restore Proxmox VE VMs directly from tapes. To restore a Proxmox VE VM, you must first restore its backups to a repository as described in the Veeam Backup & Replication User Guide, section Backup Restore from Tape to Repository.

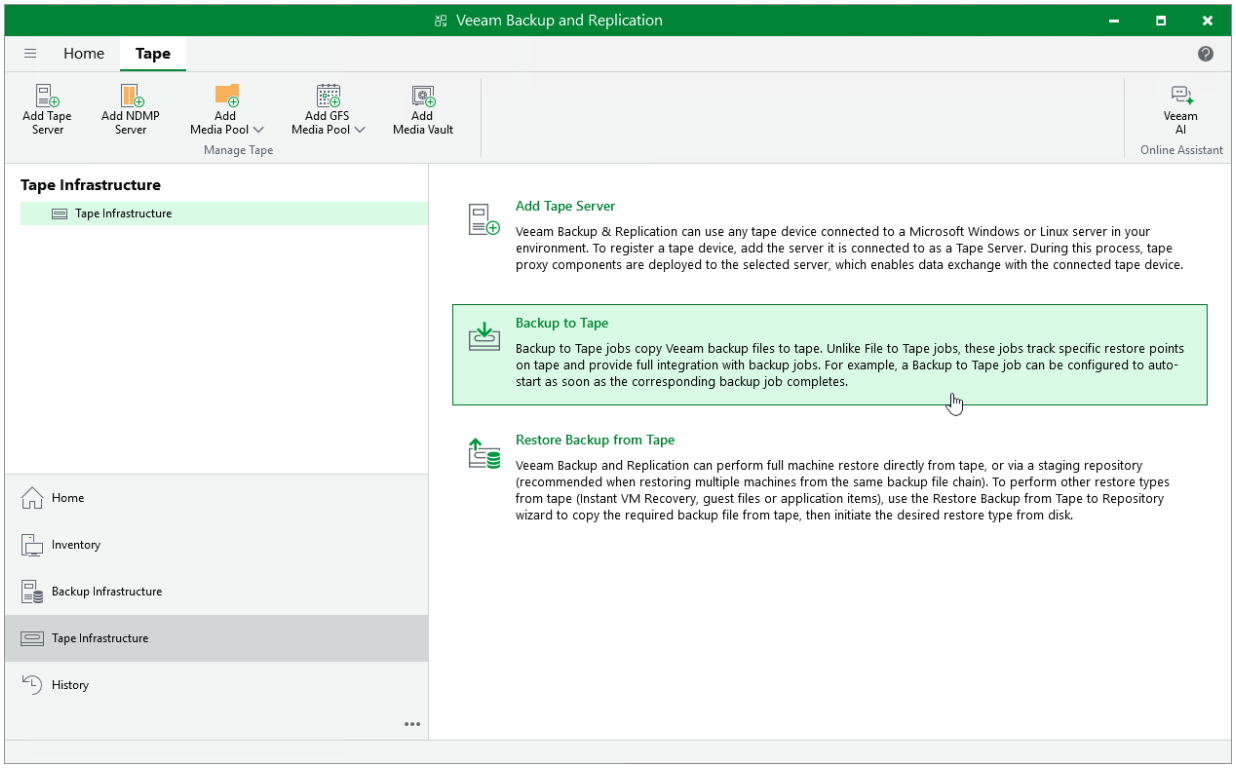

Deleting Backups

By default, Veeam Backup & Replication maintains backups stored in backup repositories according to retention policy settings saved in the backup metadata. If Veeam Backup & Replication detects that the number of restore points in the backup chain exceeds the allowed number, it automatically removes obsolete backups. You can also delete backup files from backup repositories manually if you no longer need them.

To delete backup files created for a Proxmox VE VM, do the following:

1. Open the **Home** view.
2. In the inventory pane of the **Home** view, select **Backups**.
3. In the working area, expand the job that created the backup, right-click the VM whose backups you want to delete and select **Remove from > Disk**.

   Alternatively, expand the job that created the backup, select the VM and click **Remove from** > **Disk** on the ribbon.

> **NOTE**
>
> If 4-eyes authorization is enabled in Veeam Backup & Replication, deleting backup files will require additional approval from another user with the *Veeam Backup Administrator* role.

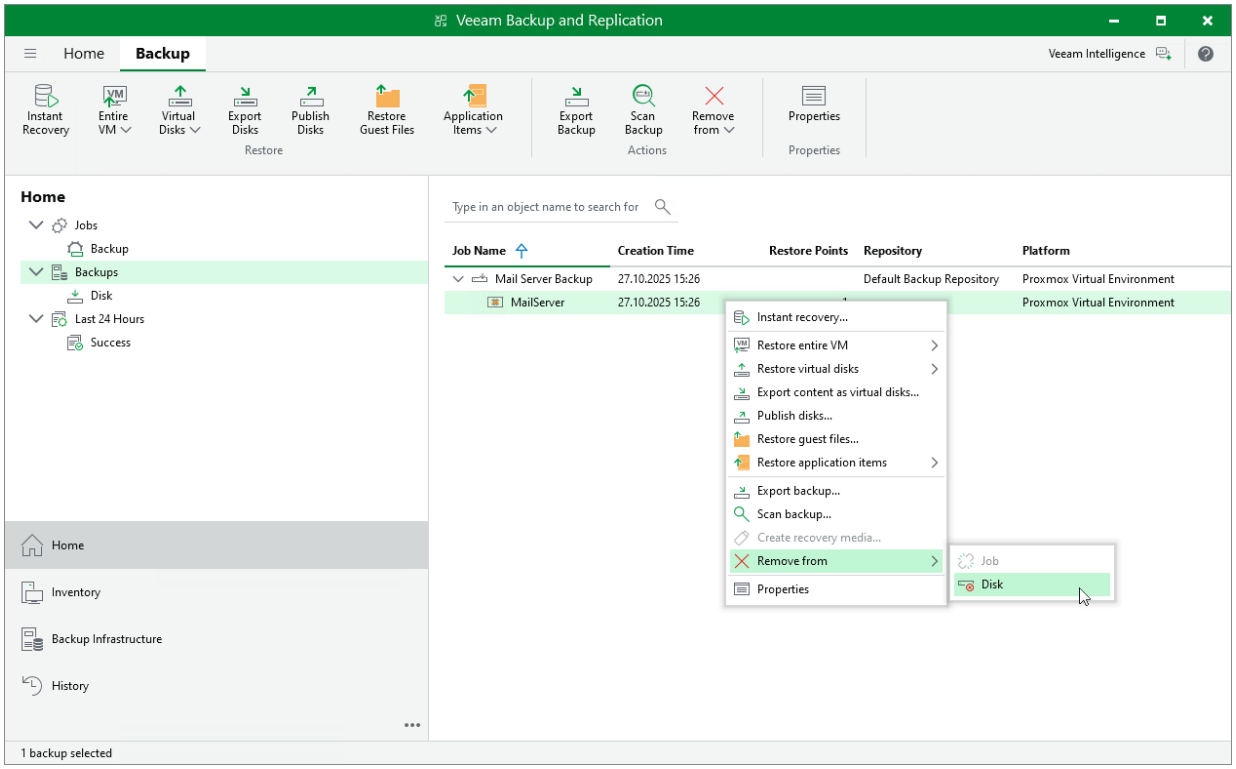

Performing Restore

In various disaster recovery scenarios, Veeam Backup & Replication allows you to perform the following operations using backed-up data:

- Entire VM restore — recover Proxmox VE VMs to the original location or to a new location.
- Instant VM recovery — instantly start an Proxmox VE VM directly from a backup.
- Disk publishing — mount specific disks of a backed-up Proxmox VE VMs to any server added to the backup infrastructure.
- File-level restore — recover individual VM guest OS files and folders.
- Application items restore — restore applications, such as Microsoft Active Directory, Microsoft Exchange, Microsoft SharePoint, and Microsoft SQL Server.
- VM disk export — restore VM disks and convert them to disks of the VMDK, VHD or VHDX format.
- VM restore to Amazon Web Services — restore Proxmox VE VMs to Amazon Web Services as EC2 instances.
- VM restore to Microsoft Azure — restore Proxmox VE VMs to Microsoft Azure as Azure VMs.
- VM restore to Google Cloud — restore Proxmox VE VMs to Google Cloud as VM instances.

You can restore VM data to the most recent state or to any available restore point.

Performing VM Restore

In case of a disaster, you can restore an entire Proxmox VE VM from a backup. Veeam Backup & Replication allows you to restore one or more VMs at a time, to the original location or to a new location.

To restore machines to Proxmox VE, you can use the following backups:

- Backups of Proxmox VE VMs created by Veeam Plug-in for Proxmox VE
- Backups of Nutanix AHV VMs created by Veeam Plug-in for Nutanix AHV
- Backups of oVirt KVM VMs created by Veeam Plug-in for Oracle Linux Virtualization Manager and Red Hat Virtualization
- Backups of Scale Computing HyperCore VMs created by Veeam Plug-in for Scale Computing HyperCore
- Backups of Microsoft Hyper-V and VMware vSphere VMs created by Veeam Backup & Replication
- Backups of VMs created by vCloud Director
- Backups of Amazon EC2 instances created by Veeam Backup for AWS
- Backups of Microsoft Azure VMs created by Veeam Backup for Microsoft Azure
- Backups of Google Cloud VM instances created by Veeam Backup for Google Cloud
- Backups of virtual and physical machines created by Veeam Agent for Microsoft Windows and Veeam Agent for Linux

VM restore is supported only for backups stored in backup repositories, object storage repositories and on the performance, capacity and archive tier of a scale-out backup repository (except for backups stored in the archive tier that consists of the Amazon S3 Glacier Instant Retrieval extent).

> **NOTE**
>
> You cannot restore VMs from backups stored in external repositories, Veeam Cloud Connect repositories, and on tapes.

To restore a protected VM, do the following:

1. Launch the Entire VM Restore wizard.
2. Select a restore point.
3. Choose a restore mode.
4. Specify a target cluster.
5. Select a storage where VM virtual disks will be stored.
6. Specify a name for the restored VM.
7. Configure network settings.
8. Specify a restore reason.
9. Verify restore settings.

Step 1. Launch Entire VM Restore Wizard

To launch the **Entire VM Restore** wizard, do the following:

1. Open the **Home** view.
2. In the inventory pane, select **Backups**.
3. In the working area, expand the necessary backup job, right-click the VM you want to restore and select **Restore entire VM** > **Proxmox VE**.

   Alternatively, expand the necessary backup job, select the VM and click **Entire VM** > **Proxmox VE** on the ribbon.

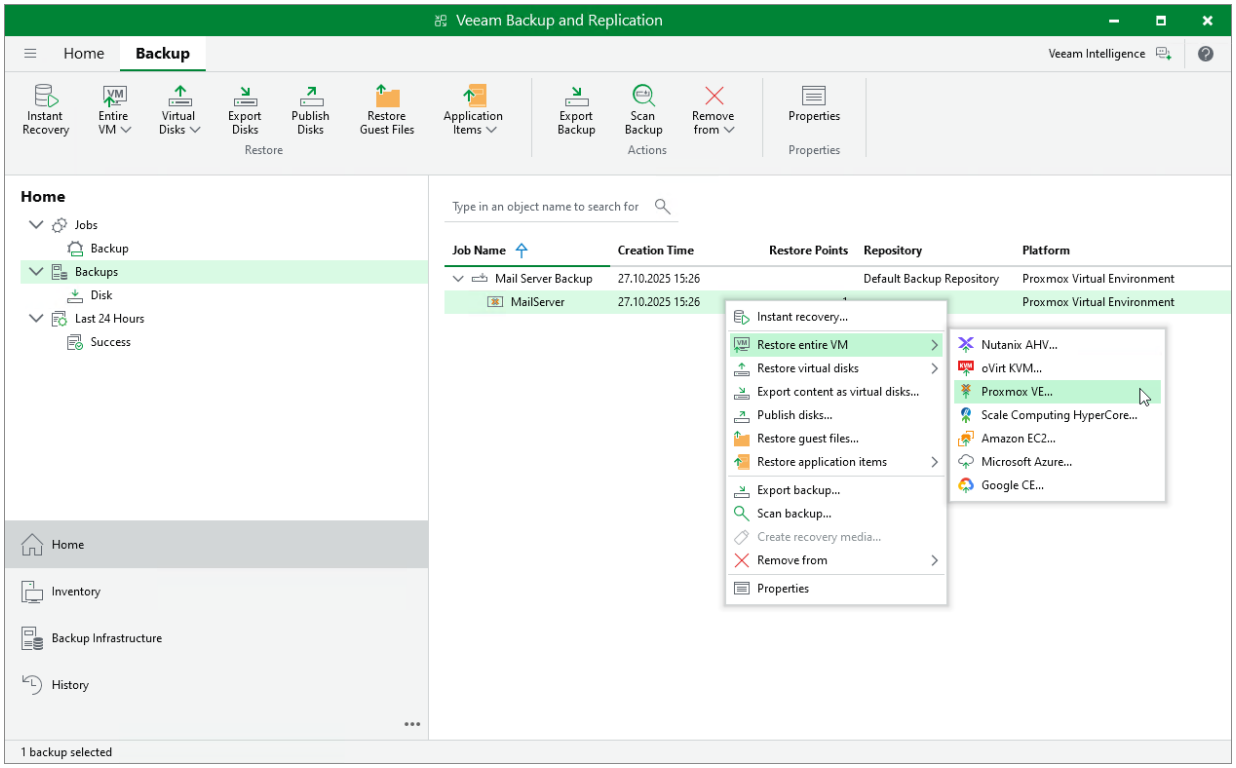

Step 2. Select Restore Point

At the **Virtual Machines** step of the wizard, select a restore point that will be used to restore the selected VM. By default, Veeam Backup & Replication uses the most recent valid restore point. However, you can restore the VM data to an earlier state.

To select a restore point, do the following:

1. Select the VM.
2. Click **Point**.
3. In the **Restore Points** window, select the necessary restore point and click **OK**.

   To help you choose a restore point, Veeam Backup & Replication provides the following information on each available restore point:

   - **Job** — the name of the backup job that created the restore point and the date when the restore point was created.
   - **Type** — the type of the restore point.
   - **Location** — the repository where the restore point is stored.

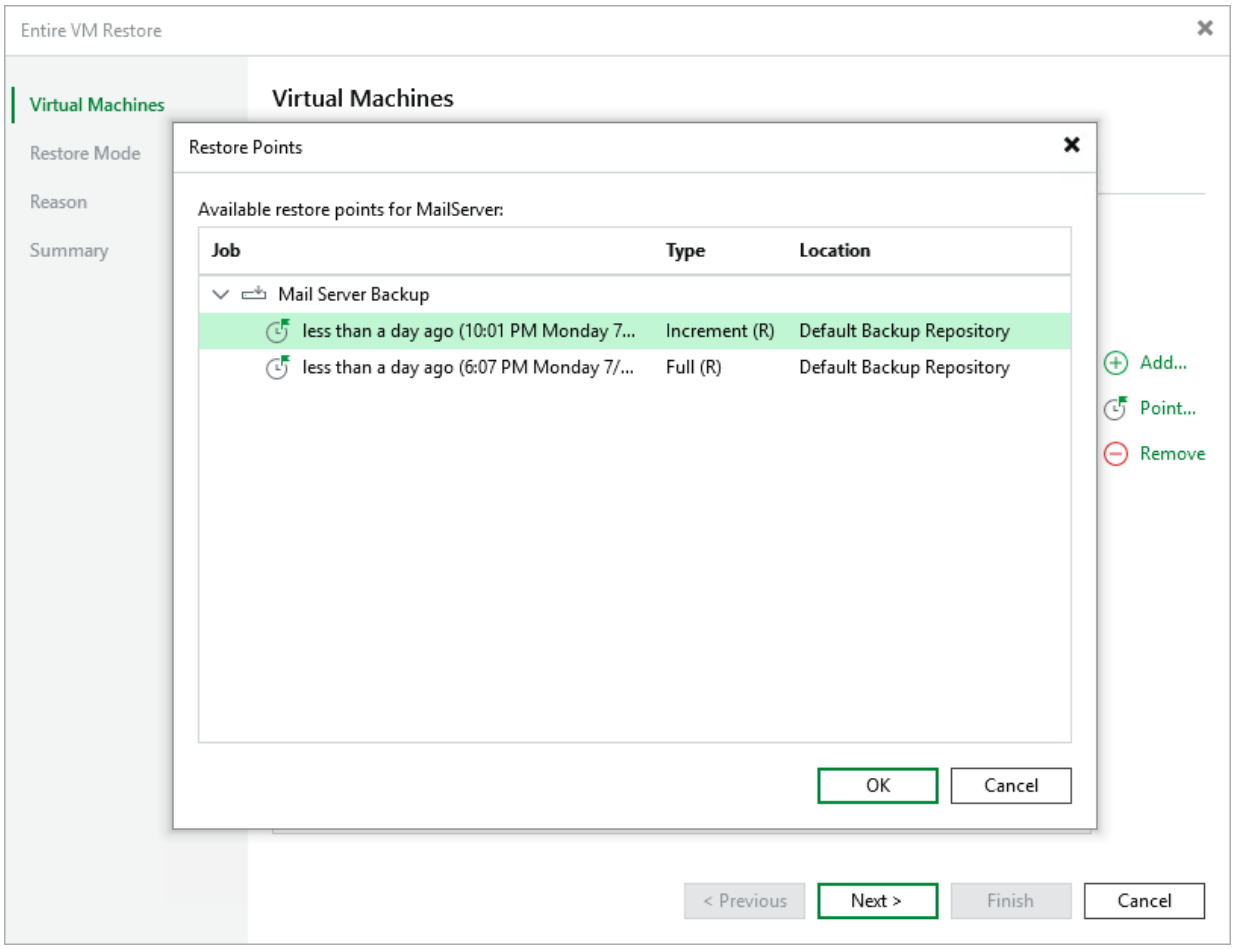

Step 3. Choose Restore Mode

At the **Restore Mode** step of the wizard, choose whether you want to restore the selected VM to the original or to a custom location.

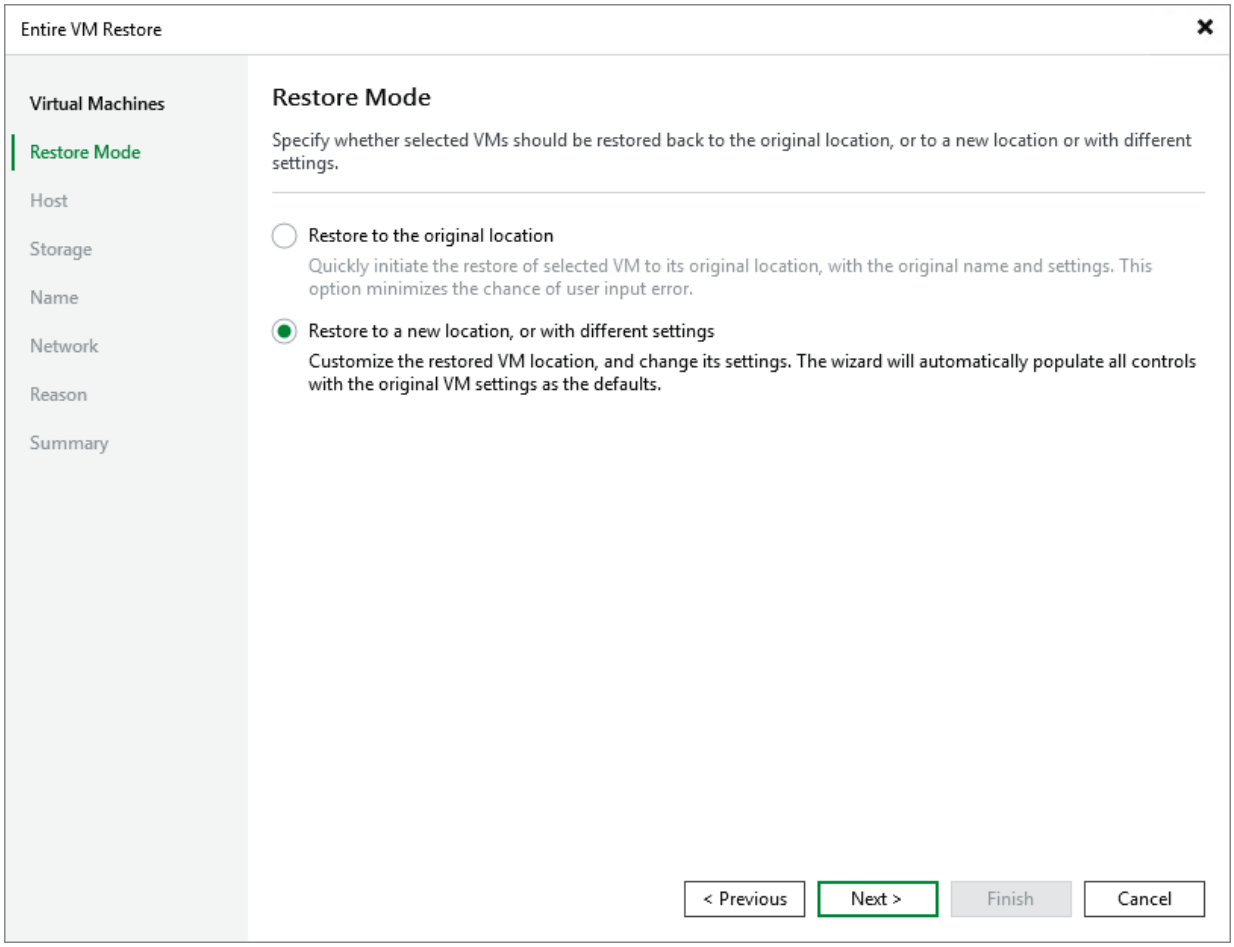

Step 4. Specify Target Host

[This step applies only if you have selected the **Restore to a new location, or with different settings** option at the **Restore Mode** step of the wizard]

At the **Host** step of the wizard, choose a host to which the recovered VM will belong. For a host to be displayed in the list of available hosts, it must be added to the backup infrastructure as described in Connecting Proxmox VE server.

> **TIP**
>
> You can also choose whether you want the recovered VM to be included into the same resource pool as the original VM.

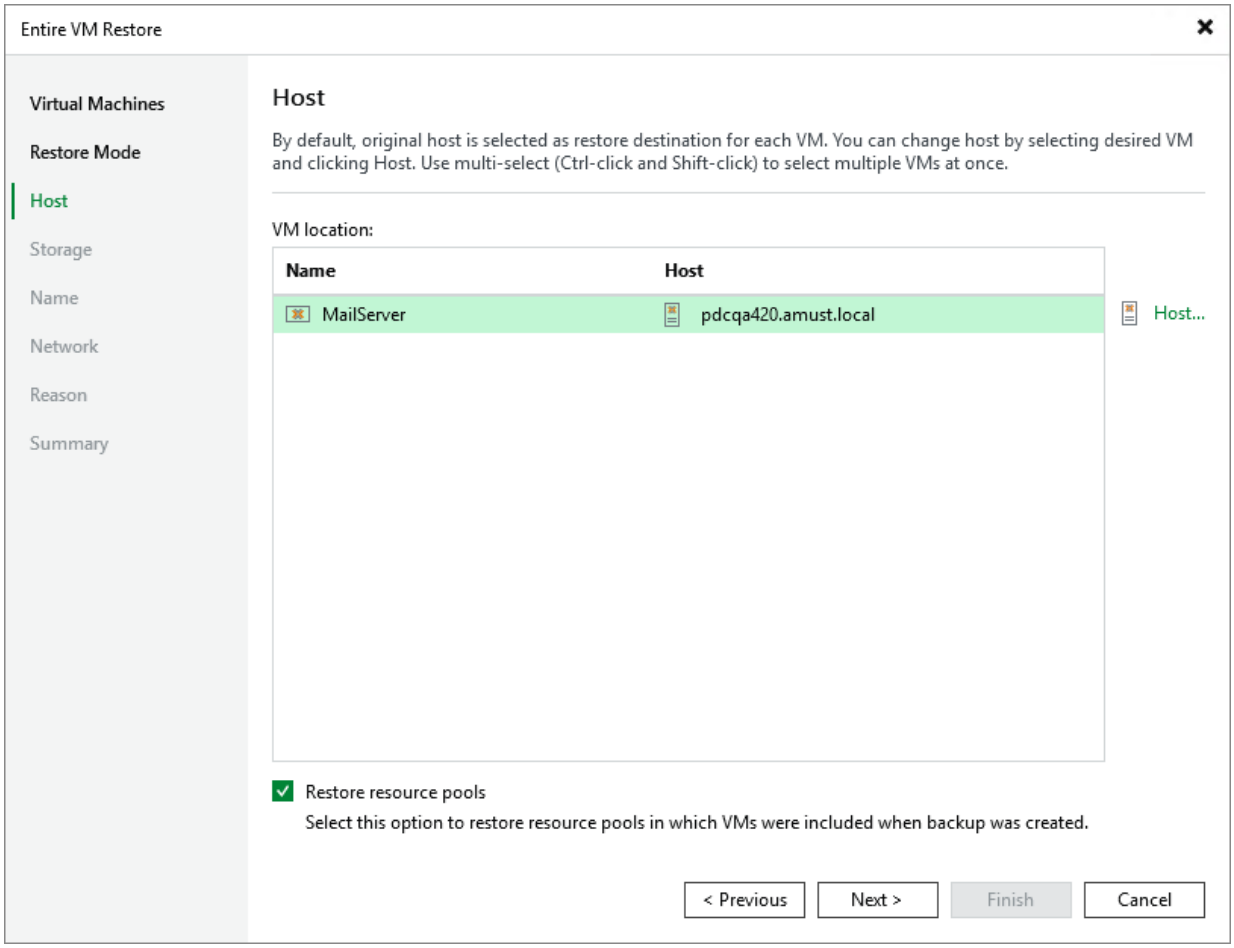

Step 5. Select Storage

[This step applies only if you have selected the **Restore to a new location, or with different settings** option at the **Restore Mode** step of the wizard]

At the **Storage** step of the wizard, choose storage where virtual disks of the recovered VM will be stored. For storage to be displayed in the list of available storage, it must be configured in the virtual environment as described in Proxmox VE documentation.

If you restore the VM to the original host, Veeam Backup & Replication will automatically select the same storage where the original VM disks were stored at the moment of backup. If you restore the VM to a new host, you will have to select storage manually. In both cases, the restored disks will by default have the same type as the original VM disks; however, you can specify another type manually.

NOTES

- You will not be able to select storage and disk type for each VM disk separately.
- If the selected storage does not support the specified disk type, Veeam Backup & Replication will display a warning notifying that some of the provided settings are invalid. You will still be able to proceed with the wizard without changing the disk type; in this case, Veeam Backup & Replication will automatically choose a supported disk type while restoring the VM.

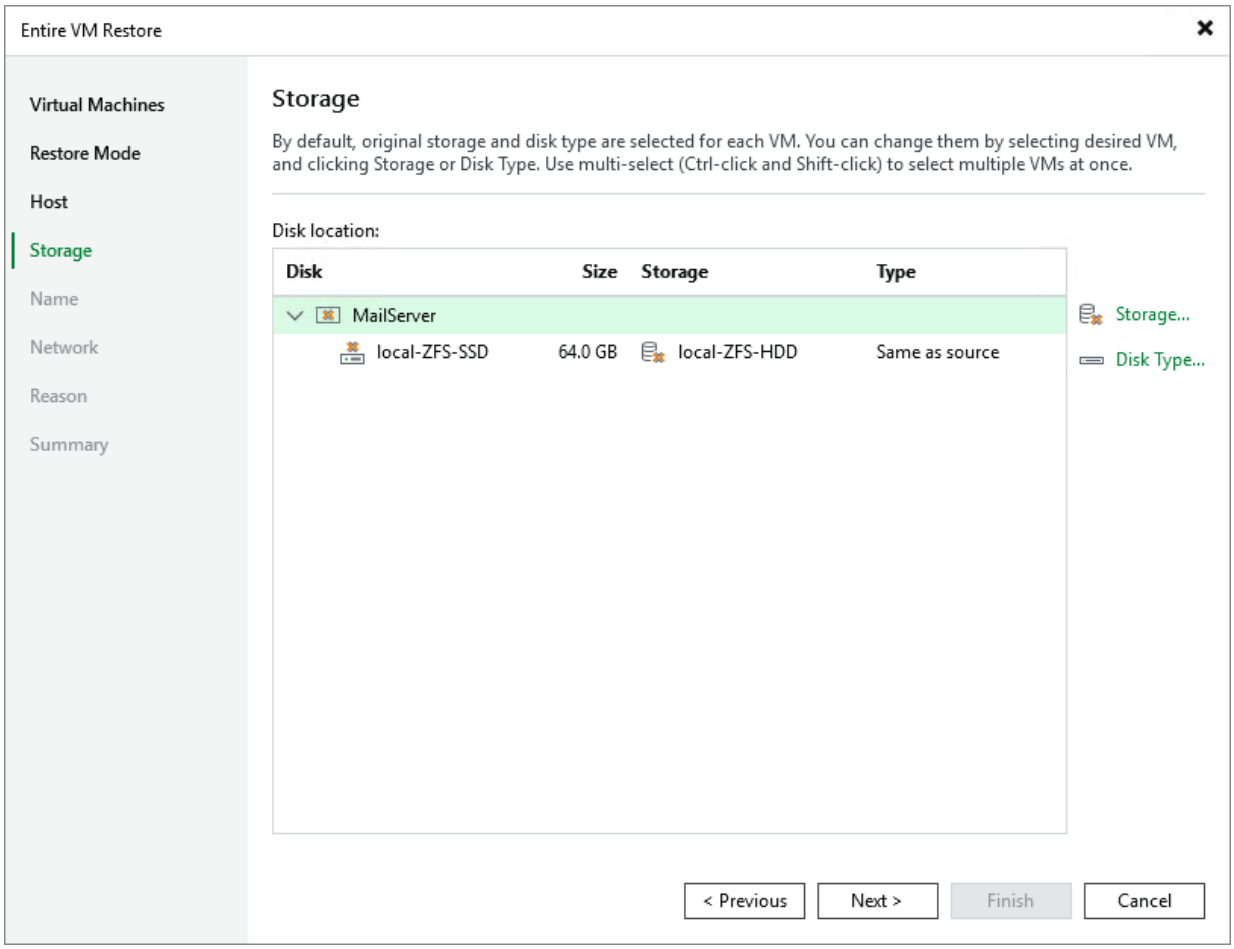

Step 6. Specify VM Name

[This step applies only if you have selected the **Restore to a new location, or with different settings** option at the **Restore Mode** step of the wizard]

At the **Name** step of the wizard, you can specify a new name for the recovered VM. The maximum length of the name is 63 characters; the following characters are only supported: a-z, A-Z, 0-9, -. The hyphen-minus character (-) is supported, but you cannot use it as the first or the last character of the name.

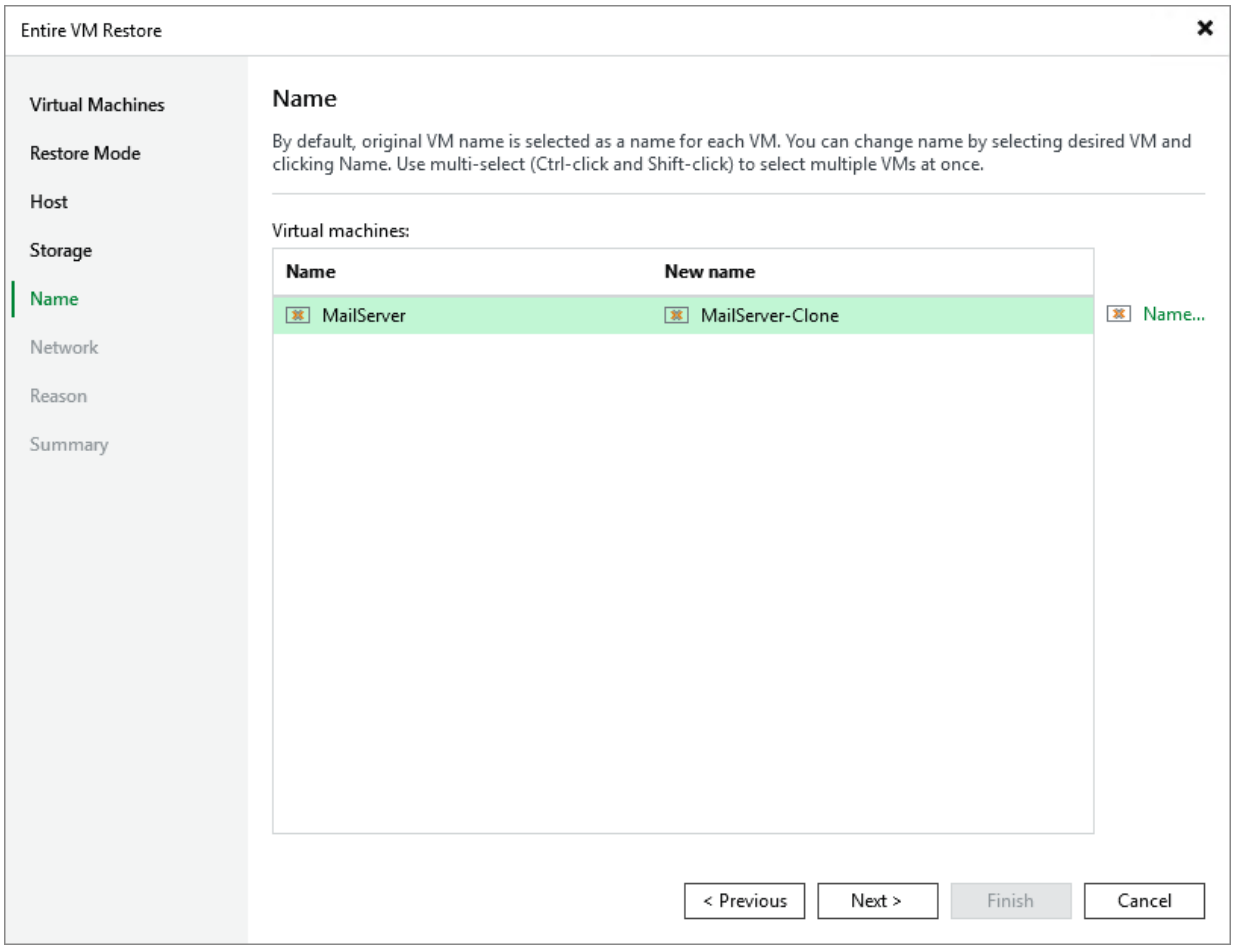

Step 7. Configure Network Settings

[This step applies only if you have selected the **Restore to a new location, or with different settings** option at the **Restore Mode** step of the wizard]

At the **Network** step of the wizard, choose a network to which the recovered VM will be connected. If you do not want to connect the VM to any virtual network, select the VM and click **Disconnect**.

For a network to be displayed in the list of available networks, it must be configured in the virtual environment as described in Proxmox VE documentation.

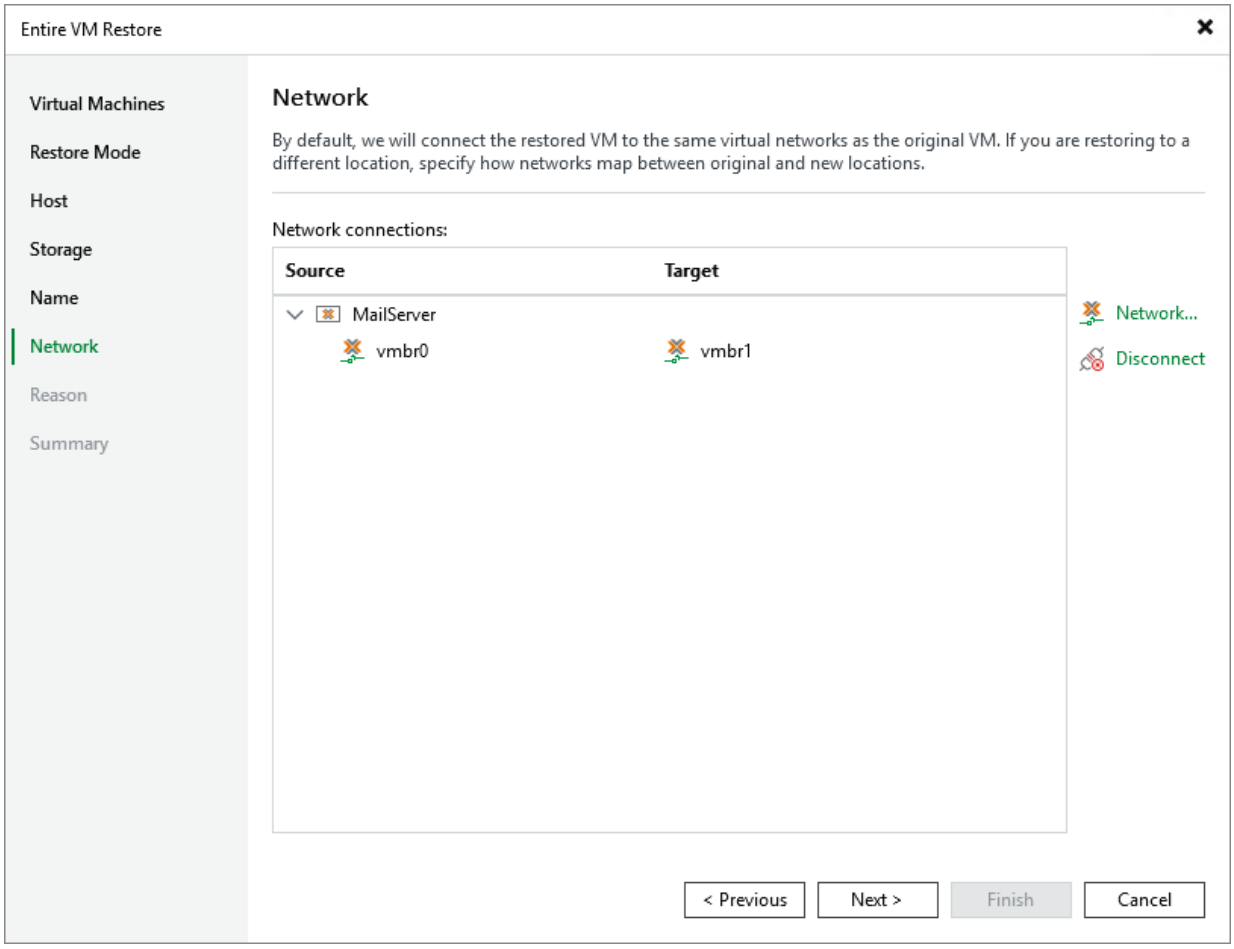

Step 8. Specify Restore Reason

At the **Reason** step of the wizard, specify a reason for restoring the VM. This information will be saved to the session history, and you will be able to reference it later.

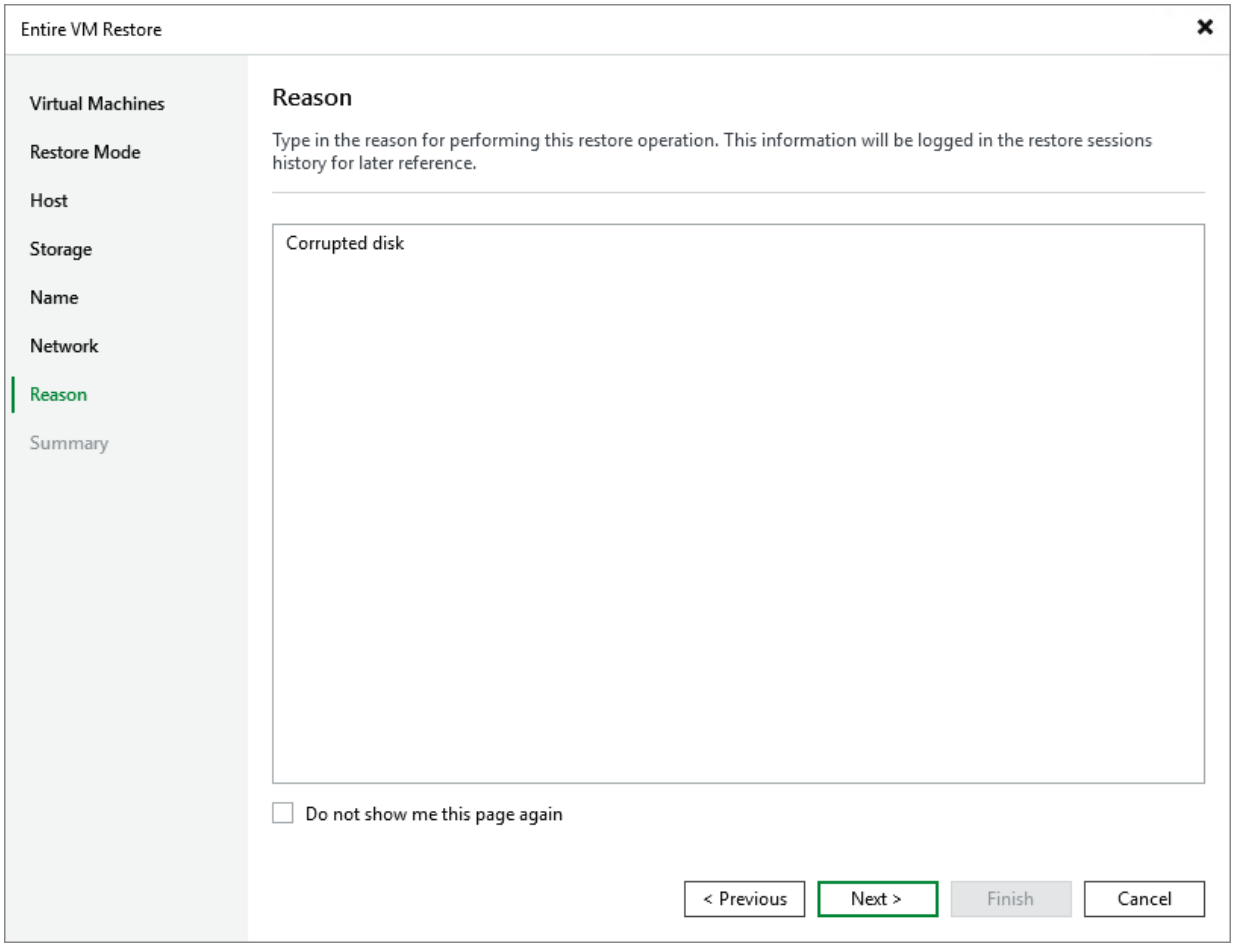

Step 9. Finish Working with Wizard

At the **Summary** step of the wizard, review summary information and click **Finish**.

> **TIP**
>
> If you want to start the recovered VM as soon as the restore process completes, select the **Power on target VM after restoring** check box.

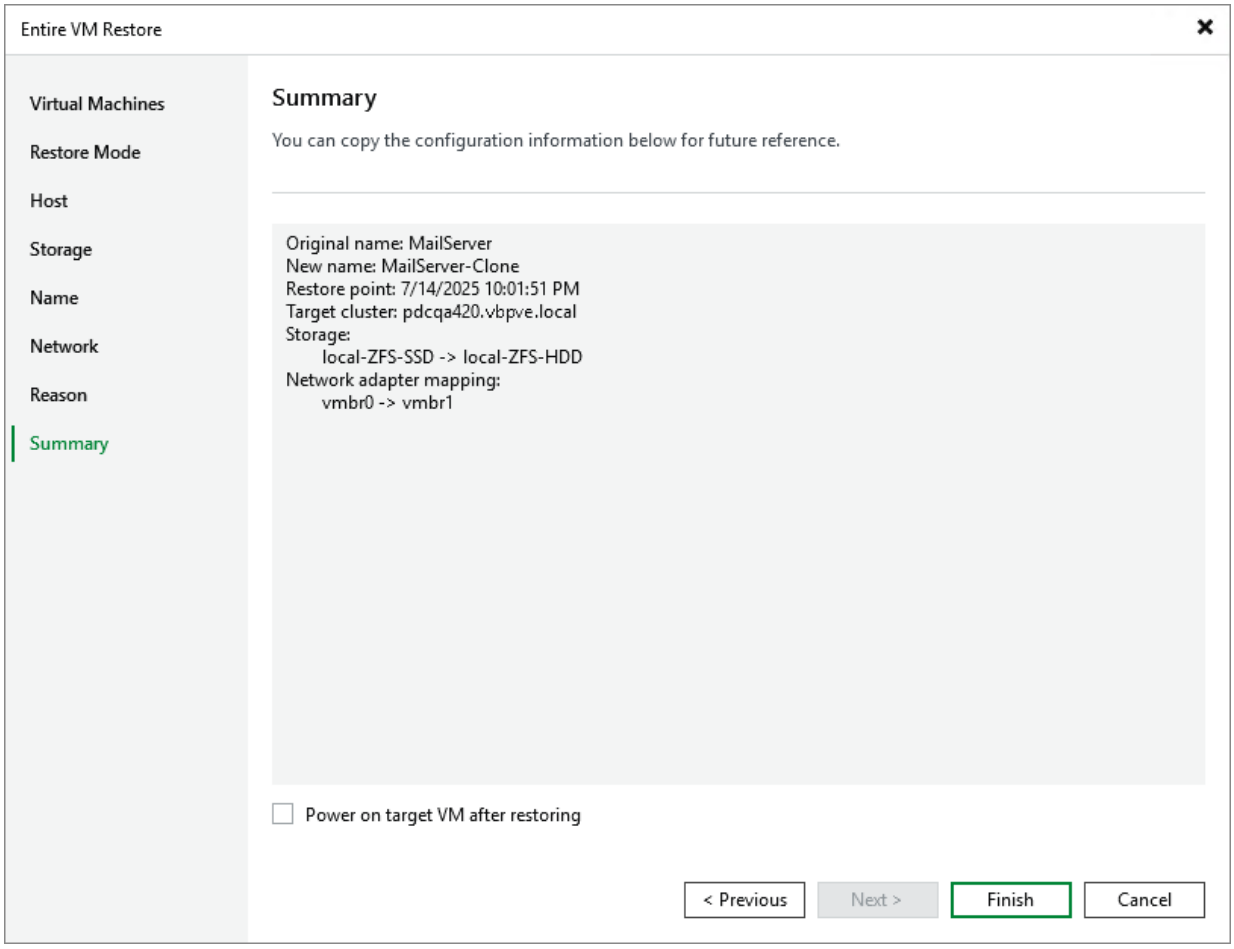

Performing Instant VM Recovery

With Instant VM Recovery, you can immediately restore Proxmox VE VMs as VMware vSphere, Microsoft Hyper-V or Nutanix AHV VMs to your production environment by running them directly from their backups. Instant VM Recovery helps you improve recovery time objectives and minimize disruption and downtime of production workloads. For more information on Instant VM Recovery, see the Veeam Backup & Replication User Guide, section VM Recovery.

To perform Instant VM Recovery, do the following:

1. Open the **Home** view.
2. In the inventory pane, select **Backups**.
3. In the working area, expand the necessary backup job, right-click the VM you want to restore and select **Instant recovery**.

   Alternatively, expand the necessary backup job, select the VM and click **Instant Recovery** on the ribbon.

   - To restore the VM to VMware vSphere, complete the **Instant Recovery** wizard as described in the Veeam Backup & Replication User Guide for VMware vSphere, section Performing Instant VM Recovery of Workloads to VMware vSphere VMs.
   - To restore the VM to Microsoft Hyper-V, complete the **Instant Recovery** wizard as described in the Veeam Backup & Replication User Guide for Microsoft Hyper-V, section Performing Instant VM Recovery of Workloads to Hyper-V VMs.
   - To restore the VM to Nutanix AHV, complete the **Instant Recovery** wizard as described in the Veeam Backup for Nutanix AHV User Guide, section Performing Instant VM Recovery of Workloads to Nutanix AHV.

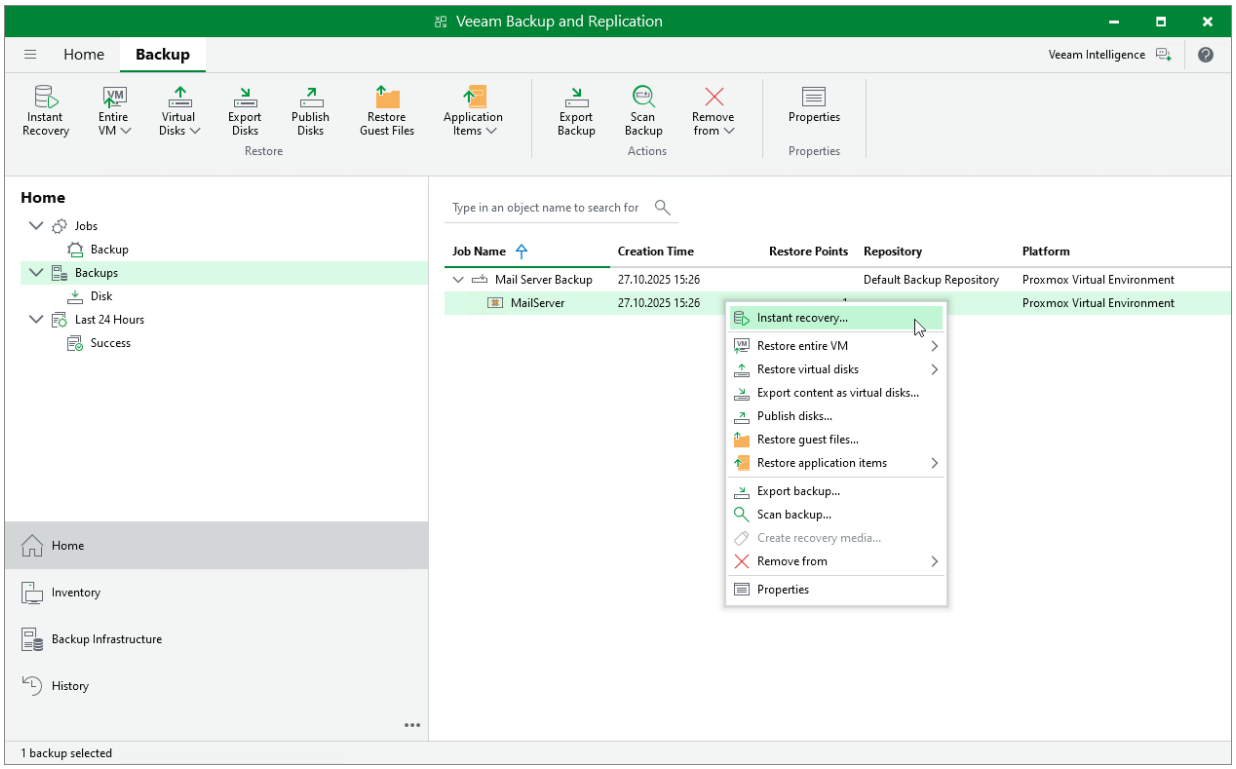

Publishing Disks

Veeam Backup & Replication allows you to mount specific disks of backed-up VMs to any server and to instantly access data in the read-only mode. This can be helpful when you want to copy files and folders as of a point-in-time state to the target server, and perform an antivirus scan of the backed-up data. For more information, see the Veeam Backup & Replication User Guide, section Disk Publishing (Data Integration API).

To publish disks of an Proxmox VE VM, do the following:

1. Open the **Home** view.
2. In the inventory pane, select **Backups**.
3. In the working area, expand the necessary backup job, right-click the VM that contains disks you want to mount and select **Publish disks**.

   Alternatively, expand the necessary backup job, select the VM and click **Publish disks** on the ribbon.
4. Complete the **Publish Disk** wizard as described in the Veeam Backup & Replication User Guide, section Publishing Disks.

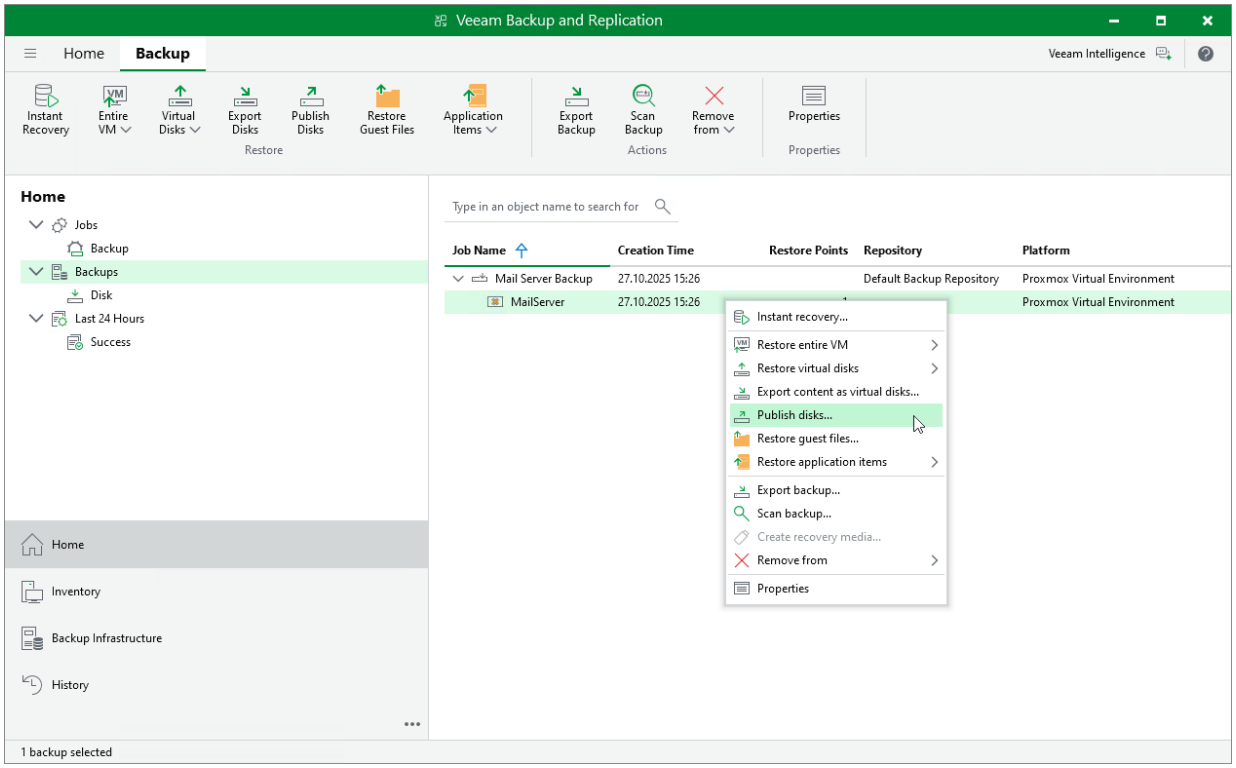

Performing File-Level Restore

With guest OS file recovery (file-level restore), you can restore individual guest OS files and folders from VM backups created with Veeam Plug-in for Proxmox VE. When restoring files and folders, you do not need to extract the VM image to a staging location or start the VM prior to restore. For more information on VM guest OS file restore, see the Veeam Backup & Replication User Guide, section Guest OS File Recovery.

> **IMPORTANT**
>
> Make sure to install the QEMU Guest Agent on VMs and enable it in the Proxmox VE administration portal — before the backups are created. You will not be able to install the agent during the recovery operation.

To restore VM guest OS files and folders, do the following:

1. Open the **Home** view.
2. In the inventory pane, select **Backups**.
3. In the working area, expand the necessary backup job, right-click the VM that contains files you want to restore and select **Restore guest files**.

   Alternatively, expand the necessary backup job, select the VM click **Restore Guest Files** on the ribbon.
4. Complete the **File Level Restore** wizard as described in the Veeam Backup & Replication User Guide, section Recovering Guest OS Files Using Console.

> **NOTE**
>
> Depending on the operating system of a VM whose files and folders you want to restore, Veeam Backup & Replication may require a mount host — a server that will be used to mount VM disks. While completing the **File Level Restore** wizard, you will be able either to choose a server already added to the backup infrastructure or to specify connection settings of a new server that will used as the mount host. For more information on how Veeam Backup & Replication selects mount hosts, see the Veeam Backup & Replication User Guide, section Mount Host Automatic Selection.

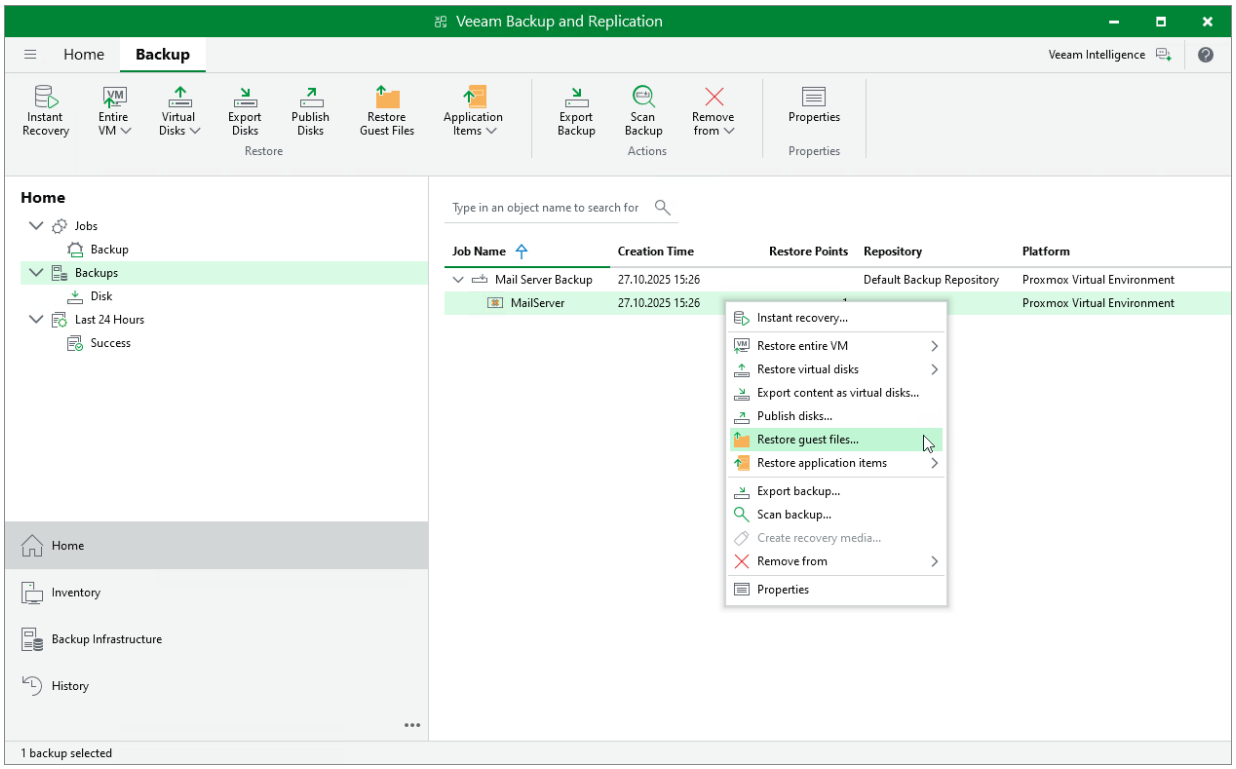
Veeam Backup and Replication
Home
Backup
Veeam Intelligence
Instant Recovery
Entire VM
Virtual Disks
Export Disks
Publish Disks
Restore Guest Files
Application Items
Restore
Export Backup
Scan Backup
Remove from
Actions
Properties
Properties
Home
Jobs
Backup
Backups
Disk
Last 24 Hours
Success
Type in an object name to search for
Job Name
Creation Time
Restore Points
Repository
Platform
Mail Server Backup
27.10.2025 15:26
Default Backup Repository
Proxmox Virtual Environment
MailServer
27.10.2025 15:26
Proxmox Virtual Environment
Instant recovery...
Restore entire VM
Restore virtual disks
Export content as virtual disks...
Publish disks...
Restore guest files...
Restore application items
Export backup...
Scan backup...
Create recovery media...
Remove from
Properties
Home
Inventory
Backup Infrastructure
History
1 backup selected

Performing Application Item Restore

With application item restore, you can use Proxmox VE backups to restore the following data:

- Microsoft Active Directory objects and containers
- Microsoft Exchange mailboxes, folders and messages
- Microsoft SharePoint sites and lists
- Microsoft SQL Server
- Oracle databases

> **NOTE**
>
> Due to technical limitations, Veeam Plug-in for Proxmox VE produces only crash-consistent (not application-consistent) backups that in some cases cannot be used for application item restore.

To restore application items from a VM backup, do the following:

1. Open the **Home** view.
2. In the inventory pane, select **Backups**.
3. In the working area, expand the necessary backup job, right-click the VM that contains an application you want to restore, select **Restore application items** and select the application.

   Alternatively, expand the necessary backup job, select the VM, click **Application Items** on the ribbon and select the application.
4. In the restore wizard, select a restore point that will be used to restore the application, specify a restore reason and click **Browse.**
5. In the Veeam Explorer application, perform the steps described in the Veeam Explorers User Guide.

> **TIP**
>
> As an alternative to application item restore, you can also perform file-level restore to recover standalone databases using Veeam Explorers.

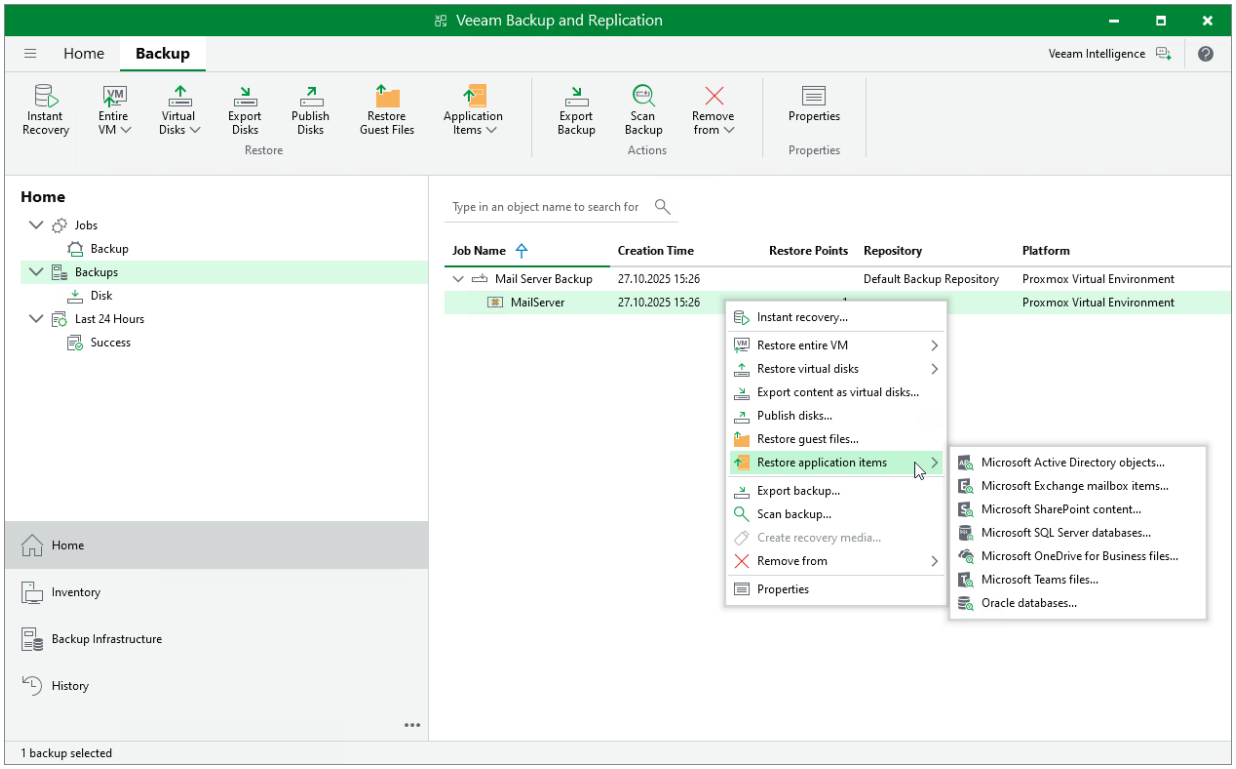
Veeam Backup and Replication
Home
Backup
Veeam Intelligence
Instant Recovery
Entire VM
Virtual Disks
Export Disks
Publish Disks
Restore Guest Files
Application Items
Restore
Export Backup
Scan Backup
Remove from
Actions
Properties
Properties
Home
Jobs
Backup
Backups
Disk
Last 24 Hours
Success
Type in an object name to search for
Job Name
Creation Time
Restore Points
Repository
Platform
Mail Server Backup
27.10.2025 15:26
Default Backup Repository
Proxmox Virtual Environment
MailServer
27.10.2025 15:26
Proxmox Virtual Environment
Instant recovery...
Restore entire VM
Restore virtual disks
Export content as virtual disks...
Publish disks...
Restore guest files...
Restore application items
Export backup...
Scan backup...
Create recovery media...
Remove from
Properties
Microsoft Active Directory objects...
Microsoft Exchange mailbox items...
Microsoft SharePoint content...
Microsoft SQL Server databases...
Microsoft OneDrive for Business files...
Microsoft Teams files...
Oracle databases...
Home
Inventory
Backup Infrastructure
History
1 backup selected

Exporting Disks

Veeam Backup & Replication allows you to export disks, that is, restore disks from VM backups and convert them to the VMDK, VHD and VHDX formats. You can save the exported disks to any server added to the backup infrastructure or place the disks on a datastore connected to an ESXi host (for the VMDK disk format only). For more information, see the Veeam Backup & Replication User Guide, section Disk Export.

To export disks of an Proxmox VE VM, do the following:

1. Open the **Home** view.
2. In the inventory pane, select **Backups**.
3. In the working area, expand the necessary backup job, right-click the VM that contains disks you want to export and select **Export content as virtual disks**.

   Alternatively, expand the necessary backup job, select the VM and click **Export Disks** on the ribbon.
4. Complete the **Export Disk** wizard as described in the Veeam Backup & Replication User Guide, section Exporting Disks.

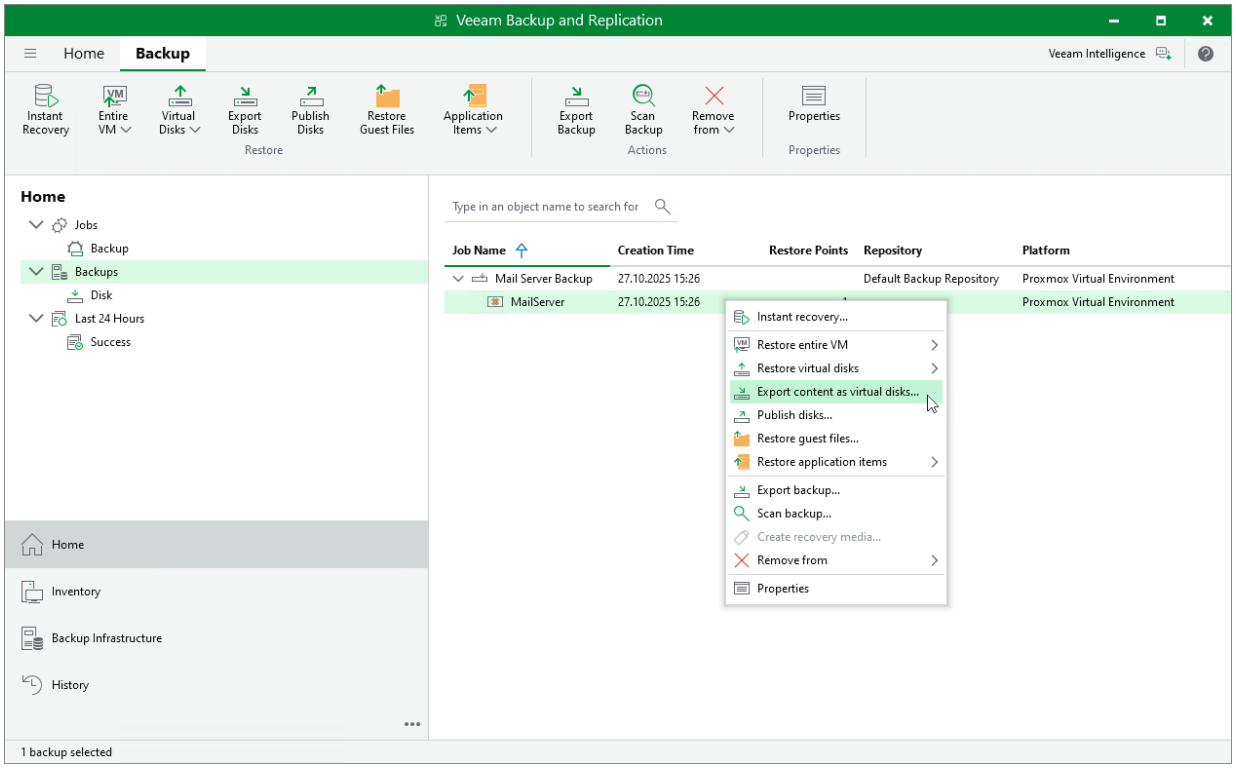

Performing VM Restore to Amazon Web Services

Veeam Backup & Replication allows you to restore Proxmox VE VMs to Amazon Web Services (AWS) as EC2 instances. For more information, see the Veeam Backup & Replication User Guide, section Restore to Amazon EC2.

To restore a VM to Amazon EC2, do the following:

1. Open the **Home** view.
2. In the inventory pane, select **Backups**.
3. In the working area, expand the necessary backup job, right-click the VM that you want to restore and select **Amazon EC2**.

   Alternatively, expand the necessary backup job, select the VM and click **Entire VM** > **Amazon EC2** on the ribbon.
4. Complete the **Restore to Amazon EC2** wizard as described in the Veeam Backup & Replication User Guide, section Restoring to Amazon EC2.

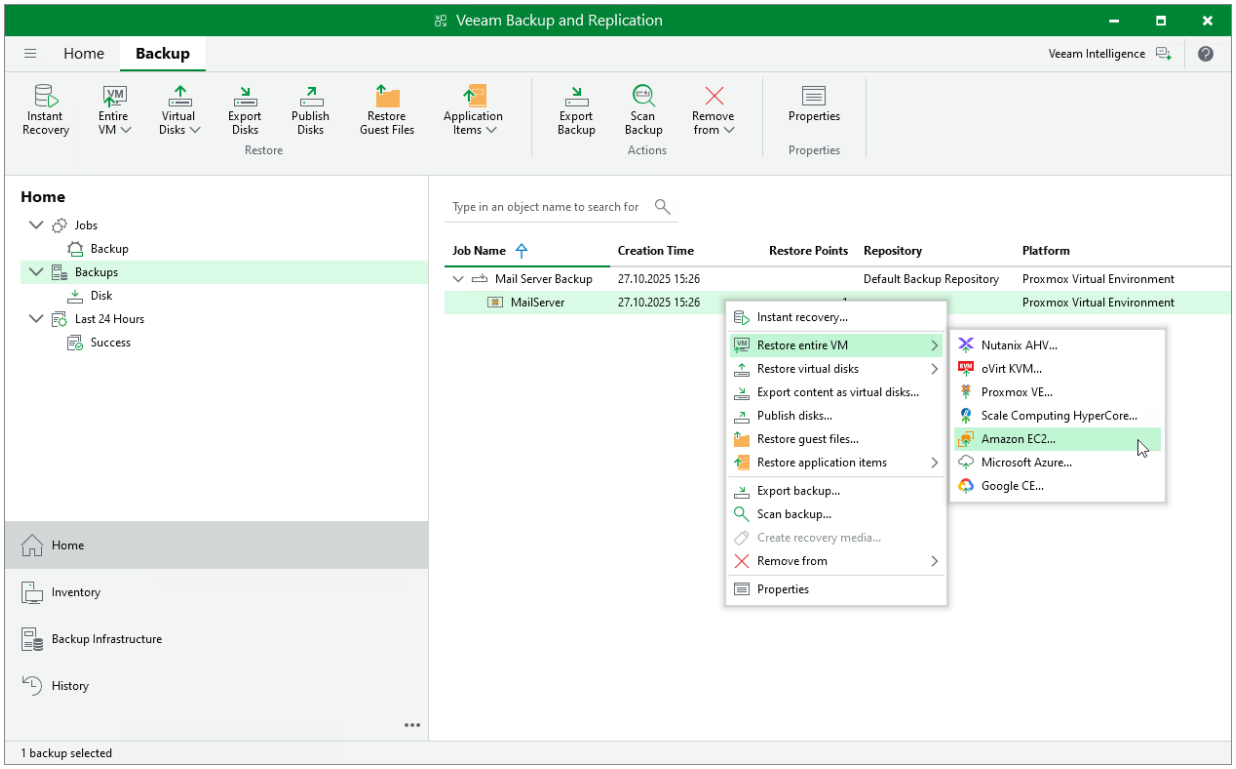

Performing VM Restore to Microsoft Azure

Veeam Backup & Replication allows you to restore Proxmox VE VMs to Microsoft Azure as Azure VMs. For more information, see the Veeam Backup & Replication User Guide, section Restore to Microsoft Azure.

To restore a VM to Microsoft Azure, do the following:

1. Open the **Home** view.
2. In the inventory pane, select **Backups**.
3. In the working area, expand the necessary backup job, right-click the VM that you want to restore and select **Microsoft Azure**.

   Alternatively, expand the necessary backup job, select the VM and click **Entire VM** > **Microsoft Azure** on the ribbon.
4. Complete the **Restore to Microsoft Azure** wizard as described in the Veeam Backup & Replication User Guide, section Restoring to Microsoft Azure.

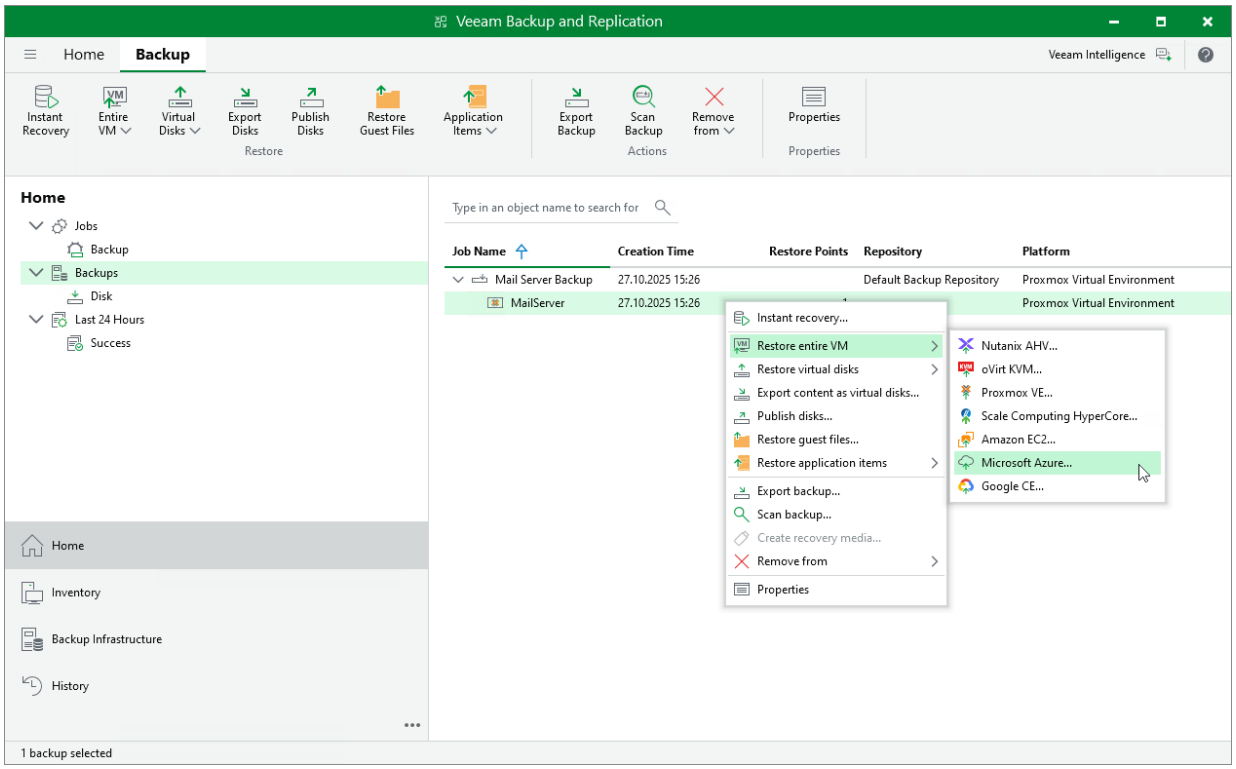

Performing VM Restore to Google Cloud

Veeam Backup & Replication allows you to restore Proxmox VE VMs to Google Cloud as VM instances. For more information, see the Veeam Backup & Replication User Guide, section Restore to Google Compute Engine.

To restore a VM to Google Cloud, do the following:

1. Open the **Home** view.
2. In the inventory pane, select **Backups**.
3. In the working area, expand the necessary backup job, right-click the VM that you want to restore and select **Google CE**.

   Alternatively, expand the necessary backup job, select the VM and click **Entire VM** > **Google CE** on the ribbon.
4. Complete the **Restore to Google Compute Engine** wizard as described in the Veeam Backup & Replication User Guide, section Restoring to Google Compute Engine.

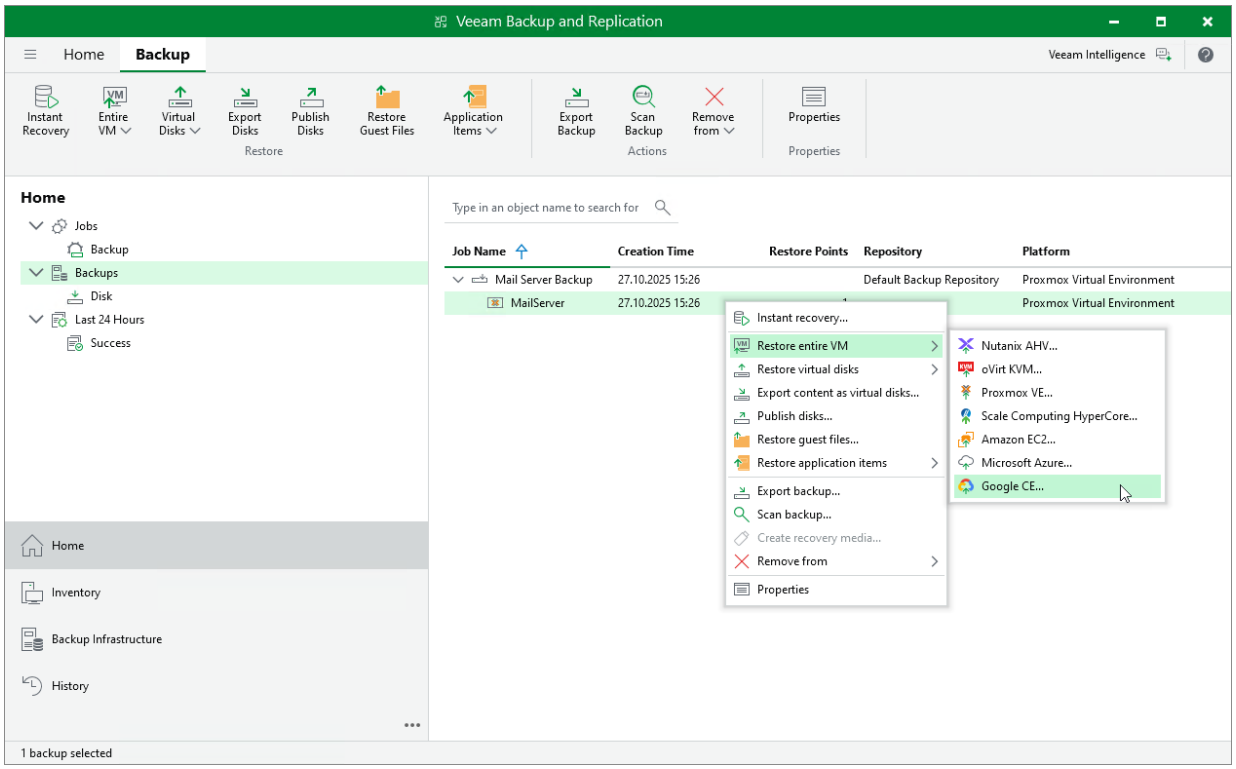

Getting Technical Support

If you have any questions or issues with Veeam Plug-in for Proxmox VE, you can search for a resolution on Veeam R&D Forums or submit a support case in the Veeam Customer Support Portal.

When you submit a support case, it is recommended that you provide the Veeam Customer Support Team with the following information:

- Version information for the product and its infrastructure components
- The error message or an accurate description of the problem you are facing
- Log files

Viewing Product Details

To view the product details, do the following:

1. On the server running the Veeam Backup & Replication console, navigate to **Settings > Apps & features**.

   Alternatively, open the **Control Panel** window and navigate to **Programs > Programs and Features**.
2. In the program list, check the version of **Veeam Plug-In for Proxmox Virtual Environment**.

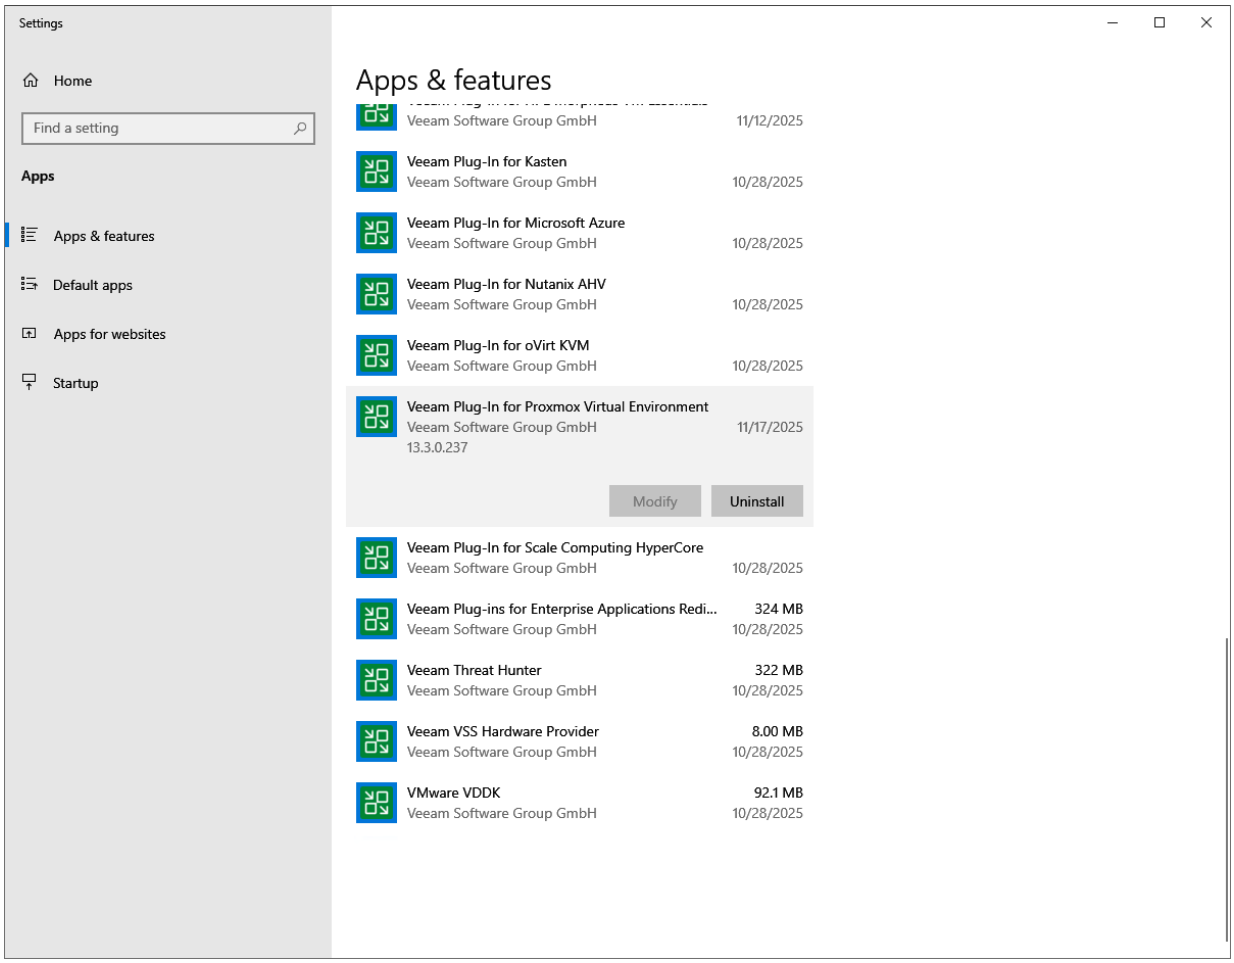

Downloading Logs

To download the product logs, do the following:

1. From the main menu of the Veeam Backup & Replication console, select **Help > Support Information**.
2. At the **Scope** step of the **Export Logs** wizard, select the **Export all logs for selected components** option. Then, in the **Managed servers** list, select the backup server.

   Complete the wizard as described in the Veeam Backup & Replication User Guide, section Exporting Logs.

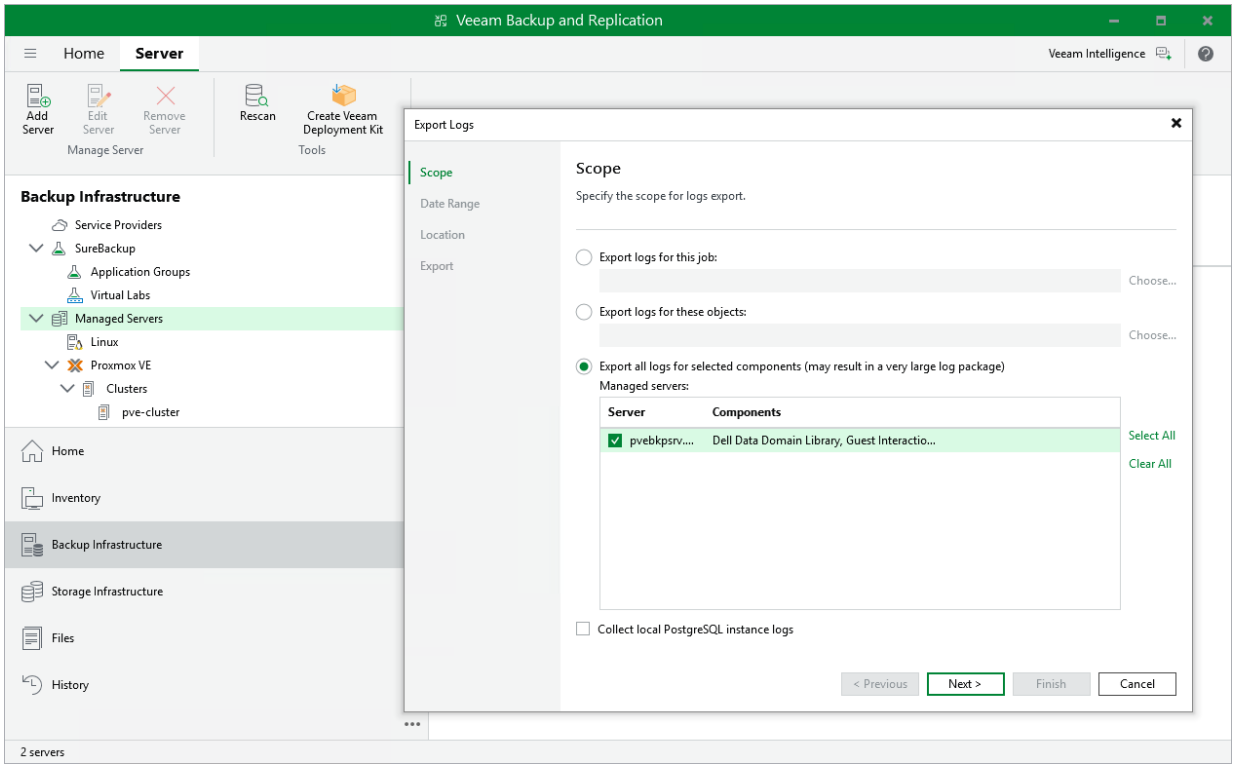

Appendix. Configuring Multiple Networks

Veeam Plug-in for Proxmox VE allows you to connect workers to multiple networks. This may be helpful if your corporate policies require that inbound and outbound internet traffic is delivered through a secure network only, or if you want to use a specific network to transfer backed-up data from and to backup repositories.

Since workers deployed by Veeam Plug-in for Proxmox VE are Linux-based VMs, they have the same limitations that apply to machines running the Rocky Linux operating system. That is, network routing can only be applied to the networks connected to the network adapters (vNICs) that have been added first while configuring workers, which means that these VMs can reach out to endpoints in other networks only through those first vNICs.

That is why you must consider the following while configuring multiple networks for workers:

- If you want workers to obtain updates from online Veeam repositories, you must connect to the first vNIC a network that allows inbound and outbound internet traffic.
- If a backup repository, the backup server or the Proxmox VE host is not reachable from the network connected to the first vNIC, you must update the worker settings to add one more vNIC and to connect it to the network to which that component is connected.

This section describes examples of valid and invalid network configurations.

Example 1. Valid Configuration

In this example, the workers, the backup server, the repository and the Proxmox VE host are connected to Network B, while the workers are also connected to Network A that allows them to obtain updates from the internet. This configuration is valid since all backup infrastructure components are connected to the same network.

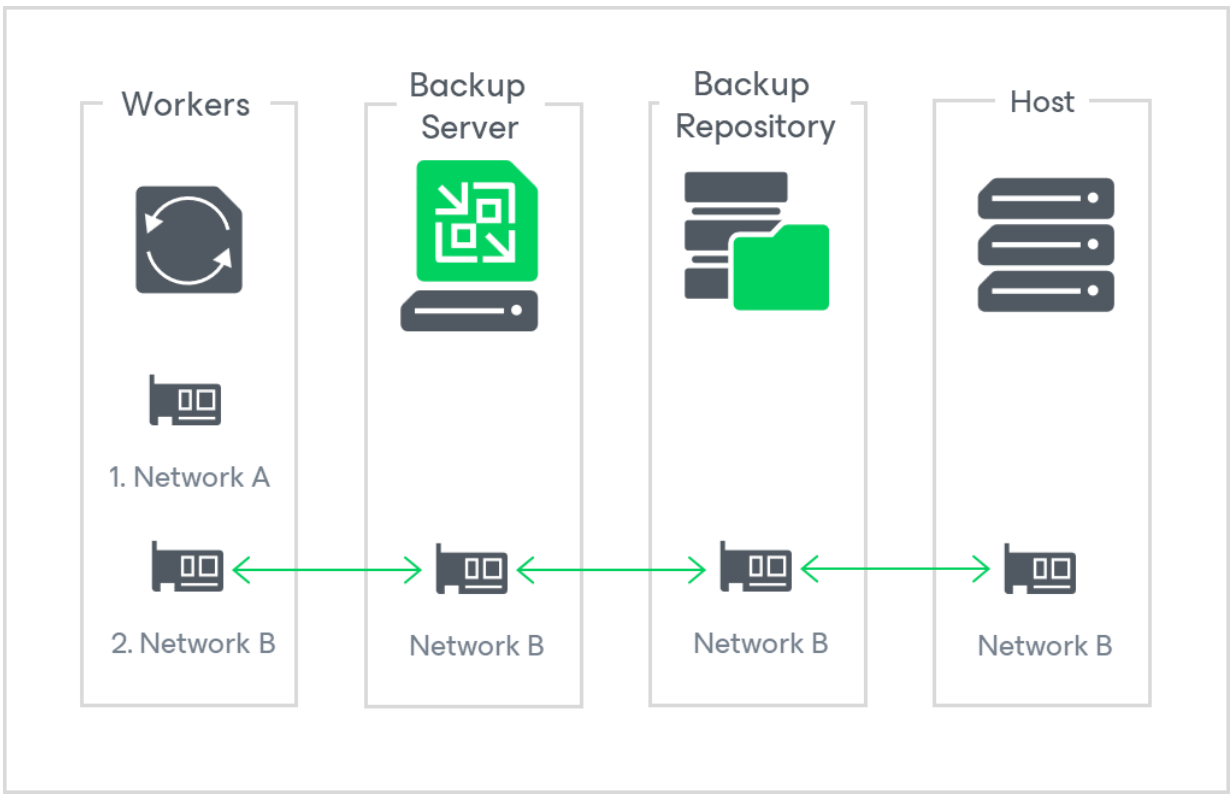

Example 2. Valid Configuration

In this example, the workers, the backup server, the repository and the Proxmox VE host are connected to Network A, while the workers and the backup repository are also connected to Network B that is configured as a preferred network to deliver traffic to the backup repository. This configuration is valid since all backup infrastructure components are connected to the same network.

NOTE

The workers will be able to obtain updates from online Veeam repositories only if Network A is configured to allow inbound and outbound internet traffic.

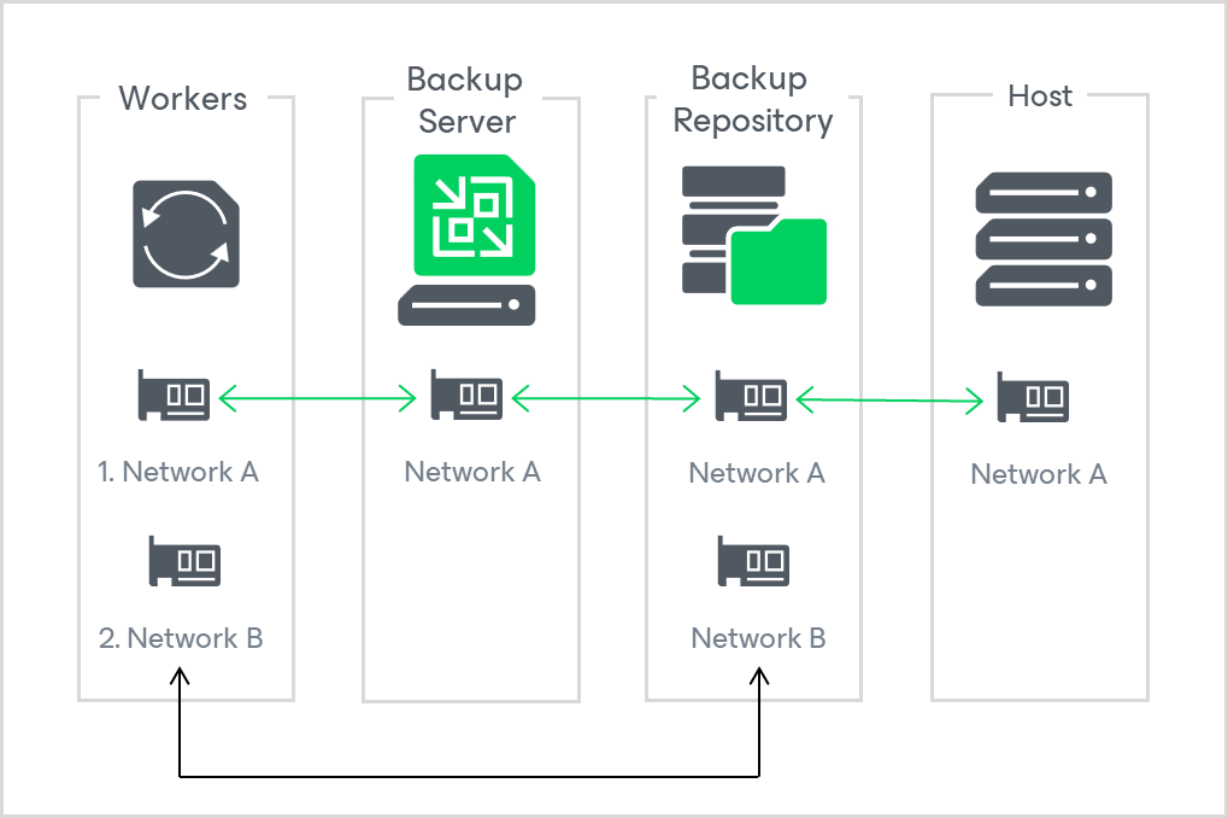

Example 3. Valid Configuration

In this example, the workers are connected to Network A using their first vNICs and to Network B that is configured as a preferred network to deliver traffic to the backup repository. Also, you have a router configured to forward traffic between networks A, C and D. This configuration is valid since the workers can use Network A to communicate with other backup infrastructure components though the router.

NOTE

The workers will be able to obtain updates from online Veeam repositories only if Network A is configured to allow inbound and outbound internet traffic.

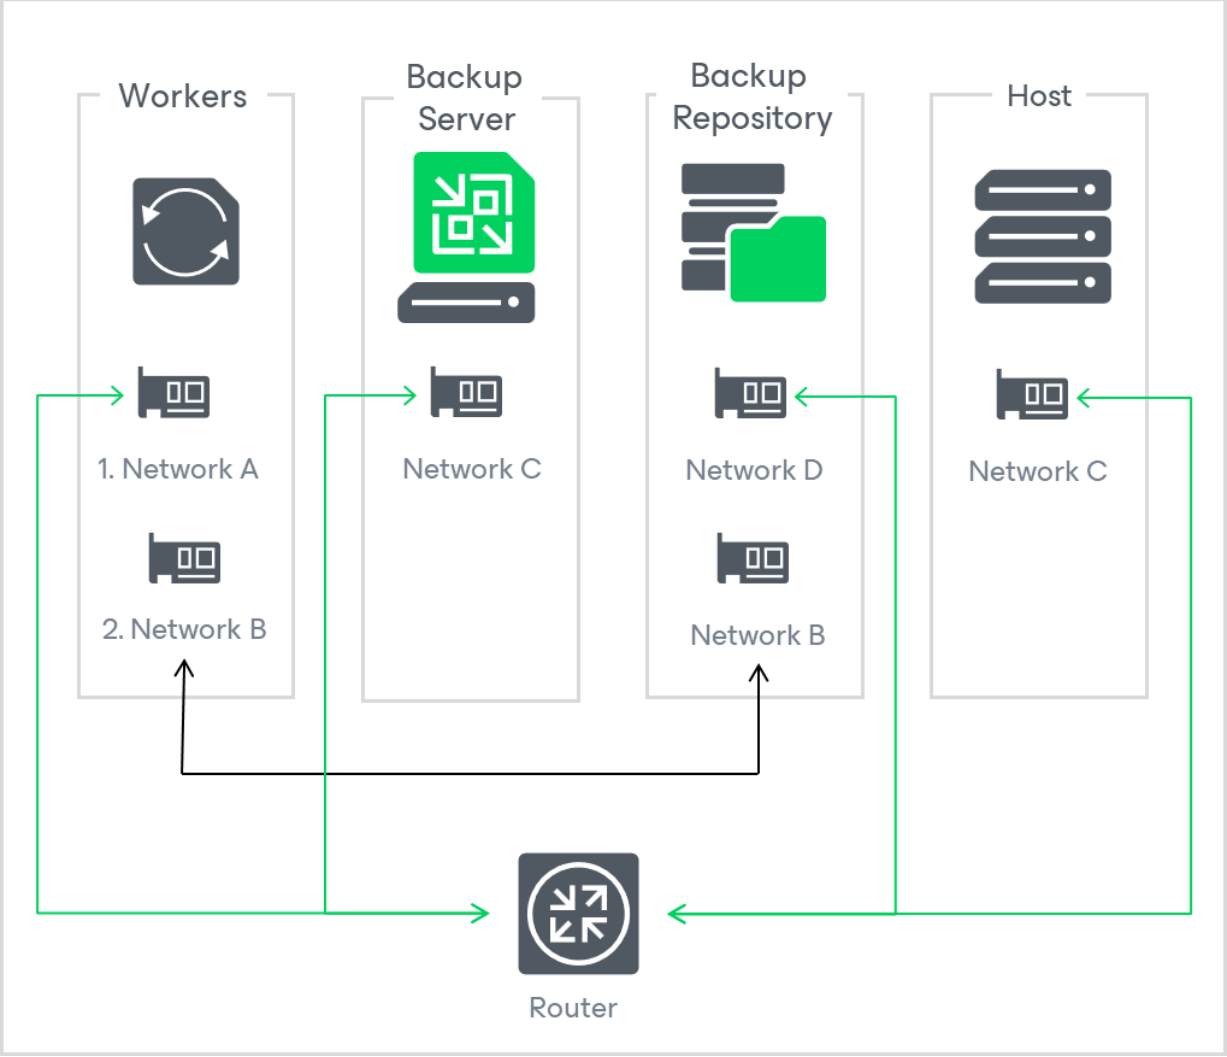

Example 4. Invalid Configuration

In this example, the workers are connected both to Network A using their first vNICs and to Network B using their second vNICs, while the backup server, the backup repository and the Proxmox VE host are connected to Network C. Also, you have a router configured to forward traffic between networks B and C. This configuration is invalid since the workers cannot use Network B to communicate with other backup infrastructure components through the router.

To make the configuration valid, do either of the following:

- Change your network configuration to connect Network A to the router.

- Add more vNICs to the workers. Then, connect these vNICs to Network C.